How business intelligence systems deliver value — Interview with Mike Beckerle of Oco, Inc.

Mike Beckerle

defines the strategic direction for and leads the development of Oco's world-class technologies. He brings 20 years of strong leadership in advanced technology with a focus on strategic large-scale information system deployments and the design, implementation and operation of scalable software and hardware systems. Most recently, Mike worked at IBM as an architect for core scalable processing capabilities for the IBM Information Server (formerly known as Ascential DataStage before it was acquired by IBM). He joined Ascential in 2001, when it acquired his startup company, Torrent Systems. As founding engineer and CTO of Torrent, Mike created the scalable computing technology that has delivered enduring value to customers from 1995 through today. Mike became an advocate of the SaaS approach in 1999, when he co-founded a start-up technology company, Fact City, which created an internet-based information service to enhance the content of high-volume internet search portals.

Keywords: *business intelligence solutions, data integration, business process management, SaaS*

Abstract This discussion covers wide-ranging issues about business intelligence (BI) systems, and how they've evolved. It explains how the limitations of current BI approaches are being overcome through new innovative approaches and Software-as-a-Service (SaaS) deployments. Oco's business model concept of "Diagonal" markets is explained, along with the synergy it has with BI delivery through SaaS to provide additional benefits.

Journal of Digital Asset Management (2008) 4, 277–290. doi:10.1057/dam.2008.22

MM: Let's start with an introduction.

MB: My name is Mike Beckerle. Let me explain a little bit about my background. I'm actually fairly new to Oco. I joined in early January of this year. As the CTO, I'm responsible for product development of new products and technologies and the strategic direction for existing products.

I joined Oco from IBM, where I was involved in large-scale computing. The core of my experience comes from spending much of my career in developing large-scale parallel processing for commercial data processing workloads. The result of this work is now the core part of the IBM information server product. I was responsible for much of the scalability of that product.

I joined Oco because it was a very good fit with my background, and my experience is very relevant to Oco's strategic direction. For example, during the dot-com era, I was involved in Software-as-a-Service ("SaaS") at a startup

called Fact City. We did something that was fundamentally SaaS (although it wasn't called that at the time) and quite similar to Oco's solution in taking data in disparate forms and using it to construct a high-volume web-accessed data service.

I decided that putting scalable commercial data processing and SaaS together would provide a great value proposition for the marketplace. I was thinking about launching this type of solution, and discovered that there was already a company doing it — Oco.

MM: Why don't you give us a quick background in terms of Oco and its core products?

MB: Oco creates "business intelligence solutions" for its customers. I use the word "solutions" because the whole business intelligence (BI) space is full of companies that sell tools, and what we do at Oco goes quite beyond tools. Let me go back briefly through the history, and then I'll come back to the point about solutions and tools.

Mike Beckerle
Oco, Inc.,
504 Totten Pond Rd.,
Waltham, MA 02451, USA
Tel: 781-810-2100
E-mail: mbeckerle@
oco-inc.com

Oco was founded in 1999. The name “Oco” actually comes from the name of the founder, a gentleman named George O’Conor. He also had the vision that many traditional BI processes were just too complex and difficult. He spent the first several years of the company’s existence developing the initial technology and testing it with a number of early customers.

Then about two years ago, the company accelerated and took in investments with the interest of ramping up growth significantly. Then of course the professional management team was brought in led by Bill Copacino as CEO.

So we have a great team, and we’re still executing on the same vision that was there originally — that is to make BI far more consumable and much easier for companies to use at much lower cost. (Figure 1)

MM: Could you give us a little bit of the history of BI?

MB: Oco was formed to address the problems of existing BI tools, which were too difficult to develop and use. I can give you the historical perspective on that.

Back in the early 1990s, people started building data warehouses, because they didn’t have access to corporate information for the purposes of reporting and data analysis. They had lots of different operational systems, but they didn’t have systems that had data from all over the place gathered together.

These projects were originally pushing the relational database technology to the breaking point. Very large data warehouses were created, and every one of the vendors struggled to make these very large databases work.

But the software has matured now, allowing companies to put together quite, quite large data warehouses. There’s now an array of companies that offer BI tools. There’s also been some consolidation in the industry lately with SAP acquiring Business Objects and IBM acquiring Cognos and so forth.

Now there’s robust relational database software out there, and there are tools for accessing the information, but it has still been much too difficult. A recent report from Gartner estimates that still over 50 per cent of these data warehousing or BI projects fail.

Figure 1: Oco technology

MM: In large part, Mike, they fail for — I think — several reasons. One — the production data sources have data that's somehow compromised or incomplete. Therefore, it requires a tremendous amount of reconditioning to make them usable, so companies can upload them into a data warehouse. Or — two — you simply have incompatible data sets from one system to another. And then I would suspect that there are probably some organizational issues around control of data, and therefore the difficulty of accessing various legacy or enterprise data sources. Does that sound about right?

MB: I would agree that the three things you've outlined are some of the factors why BI projects fail. I think that certainly there's a data quality issue, which was the first issue you raised. That's still there, although with ERP packages that issue is lessened — but certainly is not gone. Companies still have ERP systems in addition to other point-solution systems, and large companies have multiple ERP systems. So they have the problem of reconciling them. But at least within any one system, there's a sense of completeness of the information. There is still the issue of compatibility when you operate with multiple ERP systems.

There are problems that companies get into with master data management. They've got multiple ERP systems, and they have the same part or customer represented in several of them. And they don't actually have them identified in the same way, which they don't realize — and so forth. Those are significant problems in large enterprises.

This master data problem is one of the bigger challenges. At Oco, our key competency is reconciling data from multiple systems to address this problem.

MM: In fact, Mike — as we were doing the quick recap of the history of data warehouses... I think one of the things that had developed or emerged out of ERP and the data warehouses as they interact... A lot of the data that people need aren't inside the organization. They're outside the organization in suppliers and in trade partners.

MB: Well, that's certainly the case. There's information that's down in the supply chain. But in fact, in the organizations we visit the systems that are facing the supply chains at customer

sites are actually comprised of a variety of systems like warehouse management systems (WMS), transportation management systems (TMS), freight payment systems, customer relationship management systems (CRM), budgeting and financial planning systems and so forth. These systems provide functionality not found in ERP systems and therefore sit next to them — so companies always have many disparate systems where important data resides and needs to be integrated for analysis and other purposes. There still needs to be some cross-system integration of this information. ERP systems in some sense have not quite lived up to their billing of consolidating all of this information. It remains important to be able to reach across different systems as well as across multiple ERP systems to be able to provide the visibility that companies need.

You mentioned the organizational barriers. Those are really quite significant as well. I'll mention a couple of significant problems there. These are good examples for contrasting the way we approach these issues at Oco versus the way others in the industry have historically gone after these problems.

Historically, people will set out to put the data from their whole organization into the data warehouse. They're trying to get data — all the data — in one place and also must cleanse that data — an enormous task. It's an open-ended BI activity that will enable the company to utilize the data warehouse...someday.

In other words, they're building the data warehouse without already knowing exactly what they want to get out of it. They want to get whatever can come out of it.

MM: It sounds like a recipe for a very expensive digital sailboat.

MB: That's what a lot of these projects are. That's what has caused much of the difficulty and the high failure rate. There have been many successful data warehousing projects, but certainly a recipe to success is having some specific focus and purpose.

Many more benefits can accrue, but a lot of organizations simply run out of patience with the project before it has really gotten to the point where it's delivering results.

At Oco, we do something quite different. I call it the top-down approach. We basically

pick a business problem that is causing pain to the organization, and we identify a way of presenting the information to the business users in a way that we collectively believe will help them solve the problem. We create this solution by bringing our best practices and knowledge of specific functions and industries to bear. Then we work top-down from this solution design to what specific data and related information sources need to be integrated to solve that problem. So our integration work isn't open ended. We know when we are done integrating.

MM: That almost reminds me of a conversation I had with a data warehouse architect. She was building a data warehouse for an executive information system for Bank of America. She would talk about sitting down with a fairly senior marketing executive and saying, "What are the business decisions that you make in the course of a day?" And then, "What information do you need in order to make a fully-informed decision?" And "Where do you go for that information?"

Of course, there are Greenbar reports here and a conversation here and a fax here. In the course of doing that, she'd talk about identifying the most important — the number 1 or number 2 most important — business decisions that an executive would make. Then doing a map of logical but physical data sources, so as to be able to identify what the data items were that needed to be collated into information that then supported an action or an insight.

That kind of describes what you're talking about in terms of this top-down optimization strategy or top-down problem-solving sort of thing.

MB: My expectation is that a large percentage of the projects that have been successful have had practitioners working on them in the model that you just described. Here, at Oco, we've really taken that notion and turned it into an art form. We sit down with a business for one or sometimes two days and go through a systematic approach to define the key problems they need to solve. We call this approach a profiling session.

We design the solution and figure out the data resources that are going to be required and so forth. We have a quite robust methodology we go through. It's a precise recipe.

MM: Does the methodology derive from any particular established profiling methodology?

MB: No, it doesn't. It was developed internally with the input of some very experienced people who have a track record in handling complex business processes.

We do a lot of things that would probably be familiar to many database practitioners. We conduct a database dimension analysis. There are some unique aspects to our approach. The way we structure it makes it very efficient, as well as very effective, at capturing what's needed for the business people, as well as identifying the sources of the information.

MM: Is there a corresponding data diagram or an entity-relationship diagram or some other kind of high-level visual abstraction of the transformation of business data into intelligence?

MB: We actually do have a set of proprietary diagrams we use. We use a mind-mapping tool in a very powerful way, that maps the transactional data needed to solve the defined problem and all the business dimensions that would be useful in analyzing — what we call slicing and dicing — that data. And of course we bring a point of view on best practices, key metrics, and the ideal reporting and analytic frameworks that are the best ways to gain insight into a business area. We developed these with very notable experts in each functional or industry area where we work. So we bring a very thoughtful and complete starting point to the table on day one, and we work with our customers to modify these solution templates to meet specific perspectives or needs that they have. We avoid turning it into a full-up custom solution, though. So we get the best of both worlds: a world-class solution as a starting point and the tailoring of that solution to specific customer preferences and needs.

We don't publish our diagrams obviously, because they contain a lot of our intellectual property, but customers that go through the profiling process obviously get to see and benefit from that analysis.

MM: Again, we were in the middle of reprising the development of BI. You'd talked about the early days of data warehouses and then how ERP started to move through a lot of corporations, normalizing a lot of that data, giving rise to the need for a master data management as a way of harmonizing data among systems.

Then I think you were about to launch into the emergence of BI tools or technologies such

as Business Objects or Cognos or Micro Strategy or things like that.

MB: These tools, and the companies around these tools, emerged over time. There was a big flurry of tool companies that came into existence around this idea called OLAP or On-Line Analytical Processing. Its central idea was something called “Data Cubes,” which allow you to analyze and manipulate data. They give you many different ways of looking at data and organizing it along different dimensions that you need to see. You could look at items by vendor, by price or by profitability or also by geographic region, organizational roles or hierarchy, etc. The “cube” notion comes from an analogy of being able to turn a cube around in your hands to look at it from different perspectives.

These tools have been implemented in a variety of ways. In the early days, people had to summarize the data to a considerable degree in order to get these tools to perform very well. As computing power and storage has become less expensive, people have discovered that you really no longer need to summarize the data. In fact these tools become a lot more useful if you can actually drill all the way down to the lowest level of detail.

You can drill down all the way to the details and observe issues associated with the data at finer granularity. Then you are using the tool to figure out what’s causing the problem and how to solve it. This results in a much more flexible, robust and efficient solution with much faster response times.

MM: Mike — traditionally, again, I come from a background of database marketing, specifically, dealing with really large data stores. The fact is that most relational databases and most BI tools are not really good about drilling down on a what-if heuristic. Then, modeling — if I have these demographic or psychographic factors in my database, how many people does that represent?

The notion of being able to drill down into very specific sets, and almost do a little simulation in terms of the ability to access that group by e-mail, direct mail or whatever? That typically gave rise to specialized analytic databases to specifically deal with that kind of ability to drill down.

Could you just give us a quick reprise of both of those database strategies — and then get into the compare-and-contrast of them?

MB: I think the simplest way to understand it is that the relational databases were originally designed around transactional processing — the kinds of things that an ERP system needs to do. They are row-oriented. You’ll have a customer table and a customer buys something. So you put an order into the order table and there’s one row for each order, or probably each line in each order.

The whole approach is organized around the notion that there are entries — which are rows. They’re created in response to transactions. Transaction processing is the primary activity.

Then people started trying to use this to do data warehousing, where the workload is much more analytical — answering questions like “find all the customers with these characteristics” and so forth. And this required organizing the data in certain ways to support analysis and decision making versus transaction processing.

At first, people took the databases that were organized around transaction processing, and started trying to use them in different ways — to index them differently and so forth.

Then, many databases that centered around decision support entered the market. Teradata was really the first one. But there have been many since then that have entered with decision-support workloads in mind. There continue to be all kinds of interesting innovations in the database market.

The relational database market is around 30 years old. It should be mature by now, but every year there seem to be new innovations in the relational database space. I’m always astounded that there continue to be new entrants. There’s a whole slugfest among new entrants for who will have the crown of the TPC — the Transaction Processing Council benchmarks. The TPC publishes benchmarks, and they label them Benchmark A, Benchmark B, etc. They’re down to Benchmark H now. This is a decision-support benchmark. It’s really quite a good benchmark, because it measures not only how good the database is, and how fast the database can answer the question — but also how expensive it is at doing that based on a cost and a cost-performance type of metric. It’s interesting that these small vendors are continually displacing each other on the top of the heap for that benchmark, and it illustrates the continuing innovation that is occurring.

The columnar-oriented databases are databases that are organized in order to be able to handle these decision-support workloads optimally. There are quite a few of them on the market now.

But Michael, you asked the question about “what if” analysis, as well. The columnar databases and the decision-support optimized databases are very good at answering questions like “find people that have these interesting characteristics.” One needs to feed a very complex SQL query to find them, and these databases can very rapidly extract and produce that result set.

But there’s a big difference between this type of query and predictive behavior addressing questions like “If this was true, what would happen then?” These kinds of things get you into the world of data mining and predictive modeling. These types of things are — to my knowledge — not yet embedded in the databases. **MM:** Right.

Then in the history of BI, the web came along — and some things began to change. Could you quickly refresh us in terms of what changed and how as a function of the web, in the space of BI?

MB: The web changes everything. The web changes some things directly and some things indirectly. One of the interesting forces in the database world and the data processing world is that the web introduced a whole new realm of data to be handled.

The whole world of e-commerce introduced a need to understand e-commerce marketing, and to understand click-streams and how people were using the internet and so forth. That created a number of new opportunities for people to try to process and understand the wealth of data, and to understand the customer behavior.

The companies that successfully handled internet advertising have become the masters of this — Google and so forth. That’s the way that the internet raised the stakes on this kind of marketing.

There’s also the absolutely direct benefit that the web introduced — a new way to get information to people — in a way that is really much more appealing.

You’re able to get rid of many of the hassles and costs associated with software installation, if

you can just give people a website to visit to get the information they’re looking for. People really like this model. It has all of the graphical capabilities that they’ve become accustomed to with their Office and installed desktop software.

That is an immediate thing that people latch on to: “Can’t I just have this on a web page, please?” Of course there is no reason why they can’t. There are a lot of companies like Oco making that happen now.

The web also changes the way that the service, the calculations and the data preparation can all be handled. Now, and throughout the history of data warehousing — going back to the mid-1990s, there was an awful lot of outsourced data warehousing. Lots of companies outsourced their data warehousing to big companies like Acxiom that specialized in data warehouse hosting, particularly for target marketing and related applications.

The internet basically makes this idea a lot more attractive to companies — and in particular, attractive to companies with smaller budgets. It’s not just the big companies that can consider leveraging database and BI technology, but in fact, everybody now can.

People are reluctant in some cases, because they fear “Oh, gee, my precious data is going outside of my firewall.” But once people are satisfied that their data’s going to be handled securely, there are tremendous advantages.

One data-warehousing consultant I know said it pretty well: “All companies outsource the way their money is handled. That’s certainly precious to them. Why not data?”

MM: I think it’s because there’s a career track associated with it.

SaaS represents another development — almost a second- or third-wave development of the web. The idea then is that you don’t have to install software or train a whole IT service management staff to manage and provision a service. You can simply go to a provider such as Oco to get a capability that might’ve cost 5 or 10m dollars for a hundredth or thousandth of that.

MB: Yes. It really changes the game tremendously. There’s been a lot of argument over “What really is SaaS?” People have various definitions of it — some broader and some narrower. My definition of it is pretty simple.

SaaS is a service you utilize instead of buying software. It’s defined by what you don’t have to

do. You don't have to buy, learn, modify, install and maintain software.

MM: I think that the analysts have all kind of gotten together and shared some basic definitions of SaaS V1 or 1.0, which was a point solution that wasn't really set up to interoperate. It might pass data, but it wasn't really set up to interoperate with other SaaS applications or installed on-premise applications.

Then you have the second generate of SaaS, which had more of an integration platform to it, and more robust XML capability. Then you could start sharing processes among multiple SaaS applications.

MB: I think people talk about the SaaS 1.0 versus the future of SaaS. It's true that the first wave of SaaS introduced applications like Salesforce.com. Some people would even put applications like Webex into that category. I don't. The alternative to using Webex is not buying a software package. The alternative to using Webex is getting on an airplane to go give a customer presentation.

MM: I think the GoToMeeting Citrix people would probably argue with that, but that's okay.

MB: I mean the alternative to these online demo and meeting systems — Webex or the other services like it — is if you don't want to use one of those, you can't buy a package that solves the bridging problem between you and whomever you need to give a demo to. I suppose you could host such a thing on your own corporate website, but I don't recall many people doing that in the days before Webex.

In any case, the point is that these applications didn't involve integration. We have moved into an era you can call SaaS 2.0, if you want, where the applications are starting to involve the core activities or functions that businesses carry out, such as BI or ERP and so forth.

So yes, there certainly is a qualitative shift, there. But some of the industry people who I have some disagreement with would say, "It's not SaaS if you can't download it yourself" or "It's not SaaS if it doesn't have self-installation and free trial." They are basically narrowing the definition in ways that I don't believe are required. As far as I'm concerned, if an alternative to a solution requires that you have to buy software and install and maintain it, then it fits the category of SaaS.

MM: I'd like to address one issue, there, Mike. As a CTO of a SaaS company, I'm sure you'll have some things to say about this.

When I look at an enterprise application — whether it's a supply-chain management system or ERP system — generally most of these enterprise applications have anywhere from 20,000 to 50,000 function points in them. Would you concur?

MB: Yes. Very large and complex systems.

MM: Very robust. Okay.

Then when we look at the deployment of those systems, the core user of those systems barely uses 200 of the function points.

MB: Correct.

MM: Then if you look at what 95 per cent of the value is that most users generate from that system, it boils down to maybe 20–50 function points.

MB: I would agree.

MM: If you've got 50,000 function points and 50 are delivering 95 per cent of the value, what do you call the other 49,950 function points?

MB: Well, some of them are legacy — right? They're there because of the way you got to what you're selling today.

MM: What do you call it in economic terms?

MB: Low leverage. That's what I'd call it.

MM: I would call it "massive overhead."

Massive costs.

MB: Yes. But I wouldn't even blame it on the function points, interestingly enough. To me, if you have a functional aspect of your system and it's debugged and documented and so forth, then the cost of continuing to deliver that function in a new version of your product is not that high.

MM: But you know that when you come up with a new module or a new extension of it, oftentimes that's the source of the bug. A previously well-behaved, documented, debugged piece of code all of a sudden becomes the errant citizen in the new release.

MB: Well, I would say that that's true because of the...

MM: Bad architecture?

MB: No. Not necessarily even because of bad architecture. One of the reasons I work in a SaaS company is because I could no longer see a need for multi-platform packaged software any more.

MM: Exactly.

MB: Multi-packaged software has a huge issue — the combinatorics of the QA problem in the enterprise software market. The problem is that you have to support the software on every platform imaginable, and it has to connect to every version of every other package of software imaginable.

MM: And not to mention the desktop clients.

MB: Yes. The desktop clients as well. The problem here is that the combinatorics of the testing space are just huge. Function points aside — just testing the basic 50 functions across it is difficult. You end up in a space where your cost of creating and testing and delivering that software in a robust way is really expensive. The cycle time to deliver new releases is long, and the upgrade cycles become tedious and difficult. So the quality of the software is fundamentally lacking.

MM: Let me come back to my basic proposition. If 50 or 100 — it's really a talking point whether it's 50 or 200 — but when a tenth of 1 per cent of the function points deliver 95 per cent of the value, what's the value of a magic quadrant analyst thing that says, "Completeness of vision and ability to execute?" if all you have to execute on is 50 or 100 function points.

Isn't the baseline by which to compare and contrast enterprise software and/or SaaS a baseline of value in tangible, tactical, strategic or transformational? And isn't the offsetting axis — the Y-axis or counterpoint — isn't that time to value fast, moderate or slow?

So doesn't that form, essentially, the value proposition of software in general, and SaaS in particular, which is "I'm going to deliver tactical value in days or a week" and as a function of that I can deliver strategic value in days, weeks or maybe a month. And because the SaaS — as you were talking about it — the SaaS platform you can't drive innovation almost on a weekly basis... Thereby, you're driving innovation into an operation or supply chain ten, 15 or 20 times faster than these big two- or three-year build-outs on enterprise applications.

MB: I'd tend to agree. You get acceleration from the fact that as a SaaS vendor you get to select a platform and stick with it, and are able to focus on adding value for your customers. And the fact that there isn't a very difficult software deployment at the base of all your customers

means that you can roll out new functional value for them in a much more straightforward and easy way. To me, the case is so compelling.

The point you made about the magic quadrant — I'd claim that in some sense, the defenders of the magic quadrant would agree, but that all rolls up into a completeness of vision. It's a good vision if you recognize that most of those function points aren't needed.

MM: But that whole thing — it's a false proposition. The whole point of coming up with a completeness of vision and ability to execute is to sell more reports.

What utility does that have to a CIO?

MB: Basically, I think that's a dense expression of the opinions that some people have of the players in this space. But that's about it.

I want to come back to a point about SaaS deployment that is germane to where we were just going — which is security. Because once you go to these SaaS 2.0 kinds of applications, you are talking about your enterprise data being sent over to another provider that's going to give you information services of some kind. You have to place a fair amount of trust in that party.

I'll tell you an interesting thing. We — as a SaaS vendor, of course — want to be able to use our infrastructure in the most efficient and scalable way we can in providing services to those customers. That allows us to add new customers at quite a low cost, and provide a higher quality of service to them.

A lot of customers have adopted the mentality of "Well, okay. I'll let you have my data, but you'd better give me my own private server." It's because they have this sense that it will be more secure somehow. I've been trying to convince people that, in fact, it is less secure.

I tell people at this point: "Why would you want to have a private server if you're interested in security?" You're going to put it on a special, separately handled box just for you instead of putting it on the proven, secure infrastructure that we're taking care of carefully for all our customers to make sure it's secure every day.

MM: Mike, the other thing, too — if you're dealing with regulated industry — whether it's FDA or DOT or whatever... Dealing with regulated industries, as an SaaS provider, you have to meet such a high threshold in terms of transparency and IT governance — oftentimes,

higher than the user organization has in its own IT operations.

MB: Yes. That's right. This cuts both ways. If you're a large enterprise, you might be very concerned about a SaaS company being able to live up to those kinds of security standards. If you're a small enterprise or small business — an SMB business — an SaaS vendor, as an aggregator of responsibilities for data processing for multiple customers, is very likely to have a high-quality infrastructure. Higher than what an SMB company can afford to have.

There's an SaaS vendor for the credit card fraud detection space, for example. It of course has built a system to meet the highest industry standards of credit card processing for data security and so forth. It said one of the great things is that other companies feel its solution is more secure than they themselves are.

MM: Let's shift now into the conversation and hopefully extended discussion of digital supply chains and how they parallel physical supply chains to a high degree. Would you just give us a quick recap of the core concepts or the core ideas of a supply chain? Then start to correlate that to digital and physical versions?

First, supply chains start with the idea that there are multiple business entities or operations that are part of an end-to-end process of transforming raw materials or IP into some sort of tangible good or service at the end, that a consumer ends up buying.

MB: Part of the complexity of managing a supply chain is that the number of these parties is not small. If you are a company that buys things from two other companies, you do not have a big and complex supply-chain problem. But many companies buy many thousands of items from a large number of suppliers and in turn sell to hundreds or thousands of customers — these companies have very complex planning and execution issues and can benefit from new analytic tools. Beyond analytics, collaboration with your supply chain partners, sharing of information and allowing views by your suppliers into the inventory levels of their products in your facilities, or views by your customers into the status of your shipments to them can dramatically reduce your joint costs, improve product availability and increase customer service. Web-based access to these analytics and collaboration applications by

supply-chain partners is a big advantage of SaaS solutions.

MM: Isn't the other idea of a supply chain the notion of constraints or constraint theory? That is, the supply chain is as efficient as its weakest link.

MB: Well, certainly if you use the kind of lean inventory management strategies where you're trying to minimize the inventory that you're carrying — then, yes. You have to have a lot of trust that the inventory is going to be replenished rapidly by the party on the other side, and these web-based SaaS solutions provide visibility that dramatically increases this trust because each player can see what is happening, and also dramatically reduces the risk of failure.

MM: Isn't that in fact the reality of today's economy? Virtually, if you're not lean, you're carrying a whole bunch of inventory or raw materials on your income statement. So inherently you're seeing all of that as “Gee — not on my financials.”

MB: In fact, that is one of the things our solutions target. We have specific dashboards and alerts that identify for each of your many, many thousands of items, which ones are at risk of being out of stock or of having excess stock. They are solutions that allow you to run lean without getting into trouble.

MM: Isn't another dimension of supply-chain theory or supply-chain strategy, cycle time and defect rate? How quickly things move through? And how many times I have a defect or a rework at various parts?

MB: Clearly a focus on product quality, rapid processes, and rapid manufacturing and replenishment cycle times are foundational capabilities today. Companies often lack visibility, quality and cycle time metrics across the organization. These metrics are often not visible to senior management, suppliers and other key stakeholders in a company. The data may be buried in a Manufacturing Execution System (MES) or an isolated system that tracks quality, perhaps even on a spreadsheet. One of our solutions targets quality and operations reporting, and it can make these metrics or dashboards visible to all interested parties.

Consider managing inventories across the extended supply chain — from suppliers through one's own supply chain through to your customers' inventory levels. The largest players have developed

sophisticated systems to have visibility of inventory across the channel or extended supply chain. The options for the mid-sized players are much more limited and often they have not been able to afford these systems. Web-based SaaS BI solutions, such as Oco's offerings, however, now level the playing field and make these capabilities available to mid-sized companies.

MM: In your presentation at SaaScon, you had used an example of — I believe it was — Welch's, as I recall?

MB: Yes. Welch's is one of our customers.

MM: You were describing the notions of vertical BI horizontal BI. Then you coined a new term called "diagonal" BI. Could you just give us a quick recap of vertical–horizontal and then the new neologism of "diagonal?" (Figure 2)

MB: Yes. Sure. This is the business school concept — vertical and horizontal markets.

A horizontal market is a solution designed for a specific business function or application area — such as a BI software product. A horizontal market is one that can be used across industries (or across several industries).

MM: Databases. Web content management systems.

MB: That's right.

In many cases, HR packages, for example. They're not particularly industry-specialized.... They don't have any inherent industry-specific requirements built into them.

Now, in many cases, in order to use them effectively, a company that purchases one of these packages has to build in or configure in that domain knowledge or best practices themselves.

MM: In fact, they really instantiate a database or tool with a digital business model. Or at least the logic of their business model.

MB: That's right. And there is substantial cost involved in doing that.

Vertical applications involve solutions that are really specialized for particular industries. In the retail industry, you might have size assortment planning for clothing. It's absolutely specific. Not just for retailers, but for clothing retailers.

Or in the financial services area, a really good example is anti-money laundering kinds of activities. These things are very specialized for a particular industry and add tremendous value. But the number of places that you can sell such a software product is a lot smaller than one of

Figure 2: Diagonal BI

these horizontal solutions that you can sell across many industries.

MM: Would it be fair to assume that these vertical applications tend to be not necessarily transaction systems, but rather analytic systems?

MB: I do not believe that is exclusively true. But I do think more of them tend to be like that.

MM: Predominantly true?

MB: For example, in industrial manufacturing... Industrial manufacturers sell — in many cases — equipment used by other people. So they have the service and maintenance applications associated with after-market service of the equipment. That's not analytical.

Yes, they're very interested in analyzing why these machines are all failing or what the quality issues are, but they are also scheduling the repair cycles.

MM: Yes. The operational systems, then.

MB: Yes. There are some operational ones. But I would tend to agree with you that a lot of them do have the tendency to be analytical.

That brings me to the issue of what I call Diagonal BI. The term was coined by our CEO Bill Copacino. So, diagonal — if it's not horizontal and it's not vertical — what is it? Diagonal is the word we chose to define things that live across some industries, but not across all industries.

These are applications or analytical solutions that are definitely not specific to a particular industry, but they represent common functionality that's needed across many industries.

MM: In the case of Welch's, you talked about what again?

MB: That one was transportation logistics. Trucking, basically. They have to ship their goods on trucks from their plants to their customers' warehouses. It's a somewhat complex world out there in trucking. You have some customers that handle customer pick-up and drive their own truck to your factory. Other times, you have to schedule and hire trucks from a variety of carriers that deliver your products to the market.

Optimizing shipments for consumer packaged goods companies, like Welch's, can save a lot of money... A meaningful fraction of their revenue is spent on transportation costs. Companies can optimize that by assuring your trailers' utilization

is high, or by comparing accessorial charges that are the extra charges like fuel surcharges, costs for unloading of goods by the driver at customer locations, trailer storage and so forth. They can optimize these across carriers to see if any carrier's charges are out of line by analyzing shipment patterns to see where load leveling may be out of balance. By having carriers report actual time of delivery to your customers through a web-based tool, you can analyze on-time delivery performance, and enable the company to have the capabilities of larger competitors at a fraction of the cost through these SaaS web-based solutions.

MM: Not just trucks, but what's on the pallet and how many pallets get organized by what truck.

MB: That's right. And how many stops it takes and so forth.

This brings me back to what we mean by a "Diagonal" BI application.

To build an application that really helps address the problem of transportation logistics, or the truck shipping of goods, you have to embed a lot of industry understanding and knowledge of trucking into the application. So it requires information specific to the business problem of shipping goods by truck, but it's not specific to any particular industry.

You don't really care whether you're shipping machinery or consumer packaged goods or clothing. These applications cut across industries, but not all industries. Obviously, financial services people aren't shipping goods around by truck, and for the most part shipping is just not a part of their primary value proposition. Similarly, higher education is not a truck-oriented industry. But any manufacturing company, whether in the food segment, the clothing segment, the toy segment, the industrial products segment, etc, all have a similar trucking problem to solve.

Another example is any company that makes or sells something that typically has sales margin and profitability issues. The companies really want to understand what products are selling at good profit margins. They want to be assured that the inventory they carry, relative to sales rate, is in balance.

Sales margins and profitability issues cut across industries that have goods to buy and sell — but obviously these aren't applicable to

government or higher education. It's not like a database system because it doesn't apply across all industries.

These diagonal types of applications are important because they add high value for their customers. They typically save companies thousands and thousands of dollars all the time, or even millions for large companies. So they are applications that can command high price points because they really deliver great savings and a very attractive return.

But also, they're applications that, because they can be sold across many industries, have a pretty large base of prospective customers — larger than vertical-market applications that are targeting a very narrow perspective. They are very attractive from a business standpoint.

Diagonal applications also work very synergistically with SaaS deployments. That was one of the things that I emphasized in the talk I gave at SaaSCon. The reason there are companies like Oco, and obviously other new market entrants in this space, is because of this synergy.

When you build a system for a particular business problem, transportation logistics, let's say, then the structure of the database of information that's needed to support it is not specific to that particular customer. It's a database that's designed to support transportation logistics.

As a result, you can get great economy of scale in the deployment of that system by creating a SaaS multi-tenant deployment of that database. All the customers sharing that infrastructure are trying to solve the same kind of transportation and logistics problem against a database of similar structure.

This works a lot better than the ASP models of a decade ago. Back then, custom data warehouses would be designed for each business. If you tried to aggregate those together, you'd get a whole bunch of totally different databases. In some sense, they were too customized. You're not going to get common behavior by putting them together.

MM: That was one of the things that really came through in your talk, Mike. First of all, you were approaching these applications that you call "diagonal applications" really almost as a value chain optimization suite. So you're looking not just at one business, but rather at how to optimize an entire value chain — irrespective of your location in that value chain.

MB: I guess that's a way to look at it. There is a collection of these Diagonal BI applications. We've tried to package a number of them in modules that can be sold to particular industries.

MM: There was another thing that was remarkable in your presentation. That was the ability to use all of these kinds of hidden charges in the trucking area. There were some terms that you used, but they referred to — basically — "hidden" markups.

MB: Yes. That actually brings me full circle to the point I started to make at the beginning of the interview here. That was about "tools versus solutions."

When I said Oco is a provider of BI solutions, well — every business intelligence provider will tell you they're providing solutions. The question is "solutions for whom?"

If you're a data analyst, then a data analyst tool is a solution to your problem. At Oco we're trying to provide a solution to business users for a transportation cost minimization problem — like our example here. That application goes to the eyes of the business user — not to the eyes of a data analyst. It's intended for use directly by the people who are in the trenches and need that information. That's why I stress that it's a solution.

MM: The nickels and dimes, to use a metaphor. Right?

MB: Yes. Well, because it really adds up. That's the problem. This is part of the reason why summarized data cubes and so forth have given way to customers saying, "I need to be able to drill down to the actual data."

In a summarized data cube, you would just roll up all the accessorial charges noted above. If, instead, you can actually see what's happening in the individual bills of lading of the trucks, you can spot many of the problems and identify the carriers charging more than others, and so forth — even though the line-haul charge, which is the advertised cost of the shipping, is the same.

MM: I refer to these as "carbon monoxide expense items." Carbon monoxide constitutes an odorless gas that you can't see, touch or smell. But you know it is there because you have a headache. And if you're in a cave, you know the canary dies.

MB: Yes. These are, in some sense, ways for people to slide charges in on you.

MM: There was another dimension that you introduced. You kind of suggested a little bit in

terms of needing to understand the behavior of a logistics supply chain — or in this case, a transportation value chain. In classic economics, according to the work of *Ronald Coase* in his book “Theory of the Firm,” this would be referred to as “transaction costs.” Transaction costs were his way — as a theorist and economist — to describe all of the handoffs. The communication, interactions and handoffs — as well as the delays associated with getting a business process completed.

So you were really calling attention to the fact that there were all these other hidden costs — almost like opportunity costs. A percentage of the truck that wasn’t fully loaded, and the amount of time it was sitting someplace.

MB: Or the inability to ship something at a certain time, for lack of availability of capacity, and so forth.

Solving many of those problems, honestly, is easy for people once you give them access to the information.

MM: Right. Because it’s their data.

MB: Yes. It’s their data. The big headache here is integrating it from multiple systems. Representing it in a uniform way for people, getting it in the form they need and in front of the eyes of the people who have to take action on it.

In that sense, solving the transportation and logistics problem is not just a matter of some computer –science-oriented thing. It’s just as much — or more — the basics of data display and information integration.

That said, those practices have until now been far too costly and far too complex for many companies to acquire. So, that’s what we’re going after and trying to make far more cost-effective.

MM: I think it’s the function of a SaaS delivery model that allows you to have one instance of your software running in a highly secure, easily managed IT service-delivery environment. And the ability to quickly roll out new innovations — new software innovations — across your user base.

The fact that your client is basically a browser means that you don’t have a bunch of fat clients to manage and synchronize and update and all that kind of stuff.

That basically allows you to bring a strategic capability to users at a highly disruptive cost —

“Disruptive” here being the classic definition of Clayton Christiansen in “Innovator’s Dilemma” — and to see what’s next. He describes it as a “good enough solution at a significantly lower — if not almost free — cost.”

MB: Yes. What we’re doing is intended to eventually disrupt the typical development process of systems integrators building large-scale data warehousing systems and doing custom-designed data schemas and so forth.

MM: In the spirit of value-chain optimization, the idea of giving tools — not just tools, but capabilities — to the people who understand the data and who can effect an operational or a tactical decision in real time as they run their business... can you provide a couple of forward-looking comments? Or what I call “future proofs” in terms of how you see this rolling out or evolving over the next two or three years?

MB: Well, I can certainly tell you that the sophistication of these applications will be growing. I’d say over the last ten years — certainly six or seven years ago — there was a lot of talk about data mining systems, for example. But every data mining system I saw was a tool for an analyst to use.

I think this kind of advanced analytic technology is going to show up for the end users, and it’s going to be a feature of a product. It’s not going to be data mining. It’s going to show up in a way that’s meaningful to a business decision maker — as more data for them to look at and take action upon. It’s not going to be packaged as some sort of fancy data analysis tool.

I think that is certainly what our focus would be on — leveraging that kind of technology. Customers do want the capabilities that kind of technology can bring, but they don’t want it packaged for use just by data analysts.

MM: It seems to me that another dimension of your solution — the evolution of your solution — as WiFi or WiFiMax networks and/or 3G phones start to propagate more broadly... We’ve seen the great success of the iPhone and other kinds of mobile internet connected devices as being — in effect — the control mechanism. Almost like the channel changer for a television, it’s becoming the control system for these very sophisticated applications.

MB: I think of the world of mobile devices as a great way to give freedom to people who

otherwise have to be slaves to the careful tending of systems and so forth. In that sense, they're very freeing.

If you take the kind of monitoring and management application that people want as a BI solution and simply display it to them on a mobile device, you're not going to be doing them any favors. You're just changing the location at which they have to do a piece of work, where they look at a screen, make a business decision, and so forth. It might give them some location freedom, but there's a lot more potential out there for the activity you have to do — from the mobile perspective — to be at a higher level of monitoring. You automate the decision making at the lower level.

Today, let's say you're looking at a sales margin inventory kind of report. You say, "Gee. Here's a product that I have very low inventory of, and I happen to be selling a lot of it. Gee. It's selling at high margins. I guess I should reorder that."

Of course, the system should just reorder that for you.

Today, people struggle just to get all that information on one line. So they can see that the problem is actually there. The next

generation of systems will be ones directed at business rules that will help people automate the solutions. It's what we call "operational business intelligence," where triggers and tripwires and things of that sort can notice characteristics of the data in the enterprise, and can take actions.

Then, from their favorite mobile device people can make sure that the decision making that's happening for them is not going off the rails for some unforeseen reason. Instead of having to switch every switch on the train, you just have to see that the trains are all moving in a reasonable way.

I think the future will lead to integrated information properly displayed for human decision making, and to support of that human decision making.

MM: And eventually, I guess, we get into policy-managed processes that basically report back to you that "Hey. I did this. Is that okay?"

MB: Once you have integrated information, the sky is the limit with what you can do with it. Integrating the information and presenting it in a reasonable model for people has been the bottleneck and remains the bottleneck today.

MM: Well, that sounds like a great place to conclude. Thanks very much.