



BBBT Podcast Transcript



About the BBBT

The Boulder Business Intelligence Brain Trust, or BBBT, was founded in 2006 by Claudia Imhoff. Its mission is to leverage business intelligence for industry vendors, for its members, who are independent analysts and experts, and for its subscribers, who are practitioners. To accomplish this mission, the BBBT provides a variety of services, centered around vendor presentations.

For more, see: www.bbbt.us.

Vendor:	SiSense
Date recorded:	October 3, 2014
Host:	Claudia Imhoff , President, BBBT
Guest(s):	Amit Bendov , CEO
Run time:	17:35:10:00
Audio link:	Podcast
Transcript:	[See next page]



Claudia Imhoff: Hello, and welcome to this edition of the Boulder BI Brain Trust, or the BBBT. We're a gathering of international consultants, analysts, and experts in business intelligence, who meet with interesting and innovative BI companies here in beautiful Boulder, Colorado. We not only get briefed on the latest news and releases, but we share our ideas with the vendor on where the BI industry is going, and help them with their technological directions and marketing messages. I'm Claudia Imhoff and the BBBT podcasts are produced by my company, Intelligent Solutions.

I'm pleased to introduce my guest today. He is Amit Bendov. Amit is the CEO for SiSense. Welcome, Amit.

Amit Bendov: Thank you, Claudia. It's a pleasure to be here.

CI: I'm thrilled to have you here. It was a very interesting session today. Let's start off with a little bit about your company. Many people may not know who or what SiSense is. I love your tag line that you're the Robin Hood of BI, but I think it takes a little bit of explaining, so let's start with that.

AB: What we're doing, we're taking technology for understanding data that was available only to the rich folks, the Fortune 500 companies. Our vision is we'll make it available to the Fortune five million companies -- that every person in any company can make fact-based decisions without necessarily shelling out millions of dollars and spending months of struggling with data structures and all that. That's what we do. We're able to do it because of the In-chip technology. It can pack some of the processing power that used to cost millions of dollars and putting it in the hands of regular folks that don't need to worry about that technology.

CI: A little bit about the company, it started in 2004, I believe, or 2006? I can't remember.

AB: Yes, SiSense was started in 2004 by five brilliant technical guys. They're all still with the company, incredible people, technically sharp. They all come from a BI background. For six years, they pretty much sat in a crappy apartment, and worked for food, and developed this amazing technology.



Actually, I think the story is that at some point they run out of money, and they went to sit on a beach in India for a few months to develop this ElastiCube Technology. In 2010, we raised a little money from V.C.'s, not a whole lot. I hired a couple of sales guy, threw a few dollars into Google AdWords without even knowing how to sell, and sales started taking off.

I wouldn't call it in a stealth mode, but with no real aggressive marketing. We really started stepping on the gas pedal in late 2012. We grew from 30 people to 115. I always need to check because it's changing every minute.

We quadrupled the team, but sales grew tenfold. We added lots of customers. We're based in US, headquarters in New York City. We have the biggest growth area right now in the US.

CI: Growth, indeed. Let's talk about that growth a little bit. You have or you presented a buyer's profile. Who buys your product? Can you tell me a little bit about that kind of a profile that you're looking for?

AB: We usually work with people that have data challenges beyond the capabilities of Excel. They have a number of data sources that they are trying to understand together. They could be either just the variety of them, or they could be very large and want to get insights.

They could be from large companies and small companies. Again, we've nothing against the Fortune 500. We just want to make it available for everyone. Usually, for smaller companies we're their main BI solution because it's very, very quick to deploy, doesn't require technical resource or expensive hardware, and it can get insights across the company.

When in larger companies, this could be a department, maybe the marketing or sales. For example, at Dannon, they use us for finance, at Reader's Digest, for marketing, Merck for clinical vaccines.

We have some of the largest companies in the world, Target, eBay, ESPN, Kellogg's, but we're equally proud in much smaller companies. Sometimes it's two guys and a dog in that crappy apartment developing an application that requires analytics.



CI: Excellent. Let's talk a little bit about your go to marketing strategy because that is becoming critical to this organization now.

AB: Absolutely. Right now, we're 100 percent inbound, so it's all people that are coming to us.

CI: It's all word of mouth.

AB: Word of mouth and advertising. We're not trying to cold call CIOs and convince them to buy, just takes too long and not very efficient. People are coming to our website. They might read about us in Forbes or Wall Street journal or TechRanch or they click on an ad that they saw on Facebook.

They come to our website. We have a click, try, buy model. We want people to try the product because everybody can show a pretty dashboard.

The difference is evidenced when you get your hands and you try to do stuff on your own. With SiSense, you're able to do stuff very, very quickly. With other tools, often you'll struggle once you go beyond that simple spreadsheet.

CI: Let's talk about your solution a little bit. What does the package consist of?

AB: From day one, the idea was it's a complete package. It's a full stack BI solution. Everything that you need is in the box.

There is no physical box, but what you download you can get going. It has a little version of a data base or a data warehouse. We've debated whether it is or it isn't.

It is and it isn't. It's a new kind of a technology, that's the ElastiCube that collects your data from multiple sources and creates a unique model that makes it easier for people to look at it and analyze data.

It has ETL capabilities that you could extract data and transform it even if it's not clean, or you have missing values, or the field names aren't identical, or names are cryptic, you can transform it.



It has mobile and web dashboards that you could use yourself either an ad hoc or publish to your peers and your models. Everything is included, very, very easy subscription model.

You don't buy and make a lump investment up front. You pay for what you use and nothing for what you don't. You can run it on premises if you prefer or in a cloud.

CI: Excellent. Let's get back to your customers a little bit. I believe you have now over 600, which is pretty good in four years. Why don't you tell me a couple of the case studies, who's using it and how they're using it?

AB: Let's take a small company and a large company. Small company, let's look at alpharooms. It's an online reservation site within the hospitality industry. You could book hotel reservations and rental cars, flights.

They have a fairly complex data model and they wanted to understand how they're doing, how their campaigns are performing, how their finance, what people are buying and not buying, which properties or renting more or less.

They don't have technical resources to write lots of the complex sequel queries.

CI: They probably have one IT person, right?

AB: Yeah. They're using that to analyze which properties are hot, which campaigns are performing, and get the complete view of their businesses or main BI solution.

Let's take another example, maybe a large national retailer that is using us to identify shoplifting within stores and abuse of drugs within pharmacies.

They combine 10 different databases from point of sale, inventories, HR, surveillance cameras, and join 10 different databases in a fairly unique way to identify that, let's say, Gillette MACH3's are disappearing from aisle four in four stores within 25 mile radius.



This was done by a non-technical person. He's an ex-deputy sheriff by education. He doesn't know sequel. He doesn't know code, but he does know how the bad guys think.

He was able to get it done with SiSense at a fraction of the cost and time of the cost with traditional BI solutions.

CI: That's a fascinating story. It really is. He knew how they thought, and therefore he could set it up to track them. Really interesting one, I like that one.

Well, let's get into the details a little bit. What you call your secret sauce, this ElastiCube. First of all, what is it and what makes it different from a data warehouse, if you will?

AB: SiSense uses a number of technologies that are different from the traditional data warehouse. In essence, it does a very similar function. First, it's a columnar database that is not as sensitive to getting the schema exactly right or the joins or inner joins that are really way over the capabilities of mere mortals.

It is very forgiving in terms of performance. Even if you don't get it exactly right, that's OK. The schema is dynamic. You start small. You add a couple of tables. You start analyzing some data and then you add some more. Let's say now you want to add finance or HR, you do it as you go and everything is done on the fly.

The reason that we're able to do it so efficiently is twofold. First, it's the in-chip technology that leverage a lot of processing power inside a CPU.

If you look historically the traditional OLAP technology were based on disk. Then there are in-memory solutions that are faster. They use the computer's RAM, but they are limited in the size of data and performance that they can support.

SiSense uses memory inside the chip set. The -cache can store about 20 megabytes per core. It's 30 times faster than RAM. But at the same time, we're not limited by physical RAM.



We use the disk, the RAM, and the CPUs, so what you're getting is something that's faster than in mem, without the size limitation.

The other part is the user experience. If you look at the "database management" of SiSense, it doesn't look anywhere like the traditional base. The solution is very, very simple. Drag and drop. Point and click. If you're an advanced Excel user, you'll feel right at home.

CI: I think that's what's important there. The other thing that you mentioned in the presentation was that the data, the model, and the presentation are all separated from each other. I thought that was another key piece of this. Why don't you explain what you meant by that?

AB: The current data discovery tools are in vogue, started really as a desktop application that later on was server-ized. They're still facing the issues of sharing a single data model. If I say, "recognized revenue," and you say, "recognized revenue," we need to be referring to exactly the same formula of calculation and not different versions of the truth.

SiSense was built from scratch for a centralized deployment with a separation between data, modeling, and presentation. We could be looking, the data is defined as central location, so that's the relationship between a data, how they're joined together, and all the transformation definition in one place. The semantic layer is in one place.

If we wanted to do things such as "recognized revenue" and "recurrent revenue" or "age attire," this kind of formula they're shared in one place, and everybody's used the same thing. You could look at it in many ways.

You could have multiple dashboards looking at the same data model and they all use the same terminology. That's a server-based technology and not a desktop technology that is collected together.

CI: It's important to understand that difference, I think. Well, you get a couple of minutes left, but I do want to touch a little bit more on the differences or the advantages, if you will, between your in-chip capabilities versus that in-memory functionality. Why don't we just do a little compare and contrast between the two?



AB: The one thing is that you don't run out of memory. You don't get out of memory errors. The problem with in-memory is that you're confined to what can fit in that little box, whether it's 128 or 256, it doesn't matter. A lot of users spend time thinking how to fit their data, that is usually larger than a box, into that box.

That's not something that we want our users to worry about. It needs to be a seamless experience. SiSense is not limited to physical RAM. You could do a terabyte of data on a computer with only eight gigs of RAM. The other difference is the speed because we use vectorized calculation and memory inside the CPU that is faster than RAM.

It runs faster than the memory, and it scales to lots of users without getting RAM. We have one user running 400 users on an \$800 server with real fast performance.

Last, but not the least, you have a very flexible and dynamic schema without worrying about the joins.

It's really the heart that the complex data modeling. We're not saying schema free, but this is schema that you don't really need to worry about it too much. You just throw the data, connect it in some way, a little better, a little worse, but it works.

CI: I like "dynamic schema" for that very reason. There is a schema. It's just that it is flexible. We can change it without a lot of heart burn and ache and pain. I think that's again another key point in your technology that is pretty darned important.

Well, let's go into the demo a little bit. It's awfully hard to verbalize a demo, I realize that, but you did show us a number of really interesting features and capabilities. Just touch on a couple of them, a few of them.

AB: I think the interesting part about the demo is the time that it takes. My favorite trick, if I have 20 minutes, I talk for 15 and ask how much time do we have left, and I realize it's five so it's great. Now, let me show you how you get from data to insights in five minutes or less.



The speed that you can collect different data sources creates some kind of a model and get to insights is unparalleled. That said, we've seen how you could really throw data sources from multiple locations, join tables without writing code or complex sequel queries, create a short semantic model, and then start creating visualizations on top.

Share them with hundreds or thousands of users with a few clicks. Manage the security, and even we're able to see a little bit on how you can grade advance stuff with API.

You can extend the product and integrate the security model. You can create your own visualizations even if they don't exist in the product. You can integrate with your existing application with a completely different look and feel.

CI: That was a very pretty interface, I have to say. All right, you got about 30 seconds. What's the future look like?

AB: Rosy, of course. Well, we plan first to maintain our growth in terms, that's the key right now, to be less of a secret and more as a dominant player. We're focusing primarily on making sure the product works, that our support organization works. It's all the infrastructure for growth.

We're growing extremely fast, which is challenging in and of itself. Now, we need to maintain that growth, step on a gas pedal in terms of marketing and support. Plus, we have pretty radical ideas how we're going to tackle the cloud data and marry it with on prem data, again in an ElastiCube way.

CI: That's a big future for you. Unfortunately, we're out of time. That's it for this edition of the BBBT Podcast. Again, I'm Claudia Imhoff. It has been a great pleasure to speak with Amit Bendov of SiSense today. Thank you so much.

AB: My pleasure. Thank you, Claudia.

CI: I hope you enjoyed today's podcast. You'll find more podcasts from other vendors at our web site www.bbbt.us. If you want to read more about today's session, please search for our hash tag on Twitter. That's #BBBT.



And please join me again for another interview. Good bye and good business!