BBBT Podcast Transcript

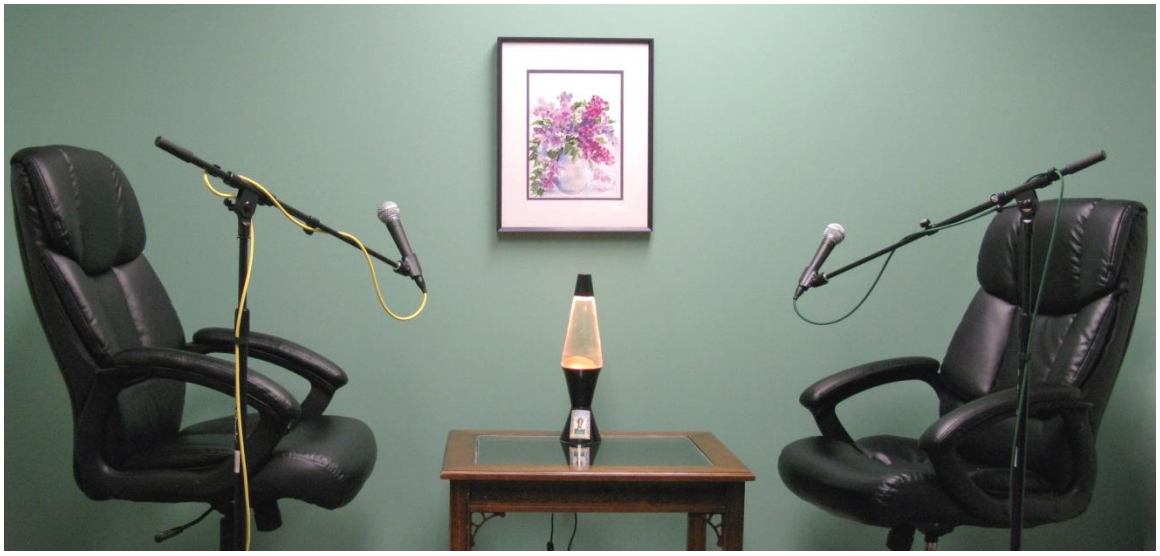

About the BBBT

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:	**EXASOL**
Date recorded:	**June 27, 2014**
Host:	**Claudia Imhoff**, President, BBBT
Guest(s):	**Aaron Auld**, CEO
	Kevin Cox, Director US Sales and Marketing
Run time:	**00:18:02**
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 guests today. They are Aaron Auld and Kevin Cox. Aaron is the CEO, and Kevin is the Director for US Sales and Marketing for EXASOL. Welcome to you both.

Aaron Auld: Thank you very much Claudia. We're pleased to be here.

Kevin Cox: Thank you very much Claudia.

CI: Alright Aaron, let's start with you. You're the CEO, so I want to hear the overview of the company. You gave a really good reason for why the company exists. So tell us about that.

AA: Yeah, I'd be happy to. You have to go all the way back to the 1990s, where we saw that the data that needed to be processed was moving further and further away from where the CPUs were and where it was being processed. So, we needed to solve that problem.

The way we did it was to look at massive parallel architecture and in-memory computing, and that has been the raison d'être of the company ever since, and that's basically the fundaments of the product we have today.

CI: Oh, that's excellent. It was an interesting problem. I mean you had to build the technology from the ground up, did you not?

AA: Yeah, we did, right from the very first line of code.

CI: There wasn't anything.

AA: No, there wasn't. I think the size the product has got to now is pretty amazing, but every single line of code has been written with in-memory processing and massive parallel computing in mind.

CI: Well, that's it. You put up a slide of the history of the innovation of the company. I thought it was remarkable. You were way ahead of your time. I mean you were doing in-memory, columnar, compression, and all kinds of things, long before many of the major relational database vendors were doing any of that, right?

AA: We were, and sometimes it's quite annoying when we see everybody jumping on that bandwagon now, and people not realizing that we have been in the game for so long. At the same time, obviously, this has been working for us as well. Because in-memory is in everybody's mind right now, so we are in the middle of a perfect storm.

We have very innovative people in the company, we have a great atmosphere for innovation, everybody's always thinking ahead, and we still do today. I think we've got some great, new innovations lined up now again today.

CI: I think one of the things that I found interesting is that you, the company EXASOL, has developed all of these capabilities. You didn't buy companies and sort of cobble them together, and say, "OK, we're now offering an analytics platform." You pretty much built it from the ground up all by yourselves. Right?

AA: We did, and if you think about it, with a very, very small team. There are only about 25 people in our development team. If you look at other companies, they've got thousands of software developers. The question is, how can you maintain quality in a team of that size? We have people who live and breathe it every day, and they've been working together for over 10 years, so it's really in our DNA.

CI: Yeah, yeah, from the ground up, like I said.

All right, well, let's talk about something that is a little bit controversial and that's the TPC H benchmarks. They were a hot item about 10 years ago,

then they kind of faded out. Now they're coming back again. But EXASOL has been knocking the ball out of the park on these TPC H benchmarks.

First of all, talk about your performance in that benchmark. And I think the second thing is do you see them as relevant today?

AA: Our performance has been pretty amazing, which is what you get with an in-memory technology. We're so far ahead of the competition that sometimes we've thought that the benchmark is actually dead.

On the other hand, and this talks to your question about the relevance of the benchmark, it is the only benchmark out there. As long as nobody comes up with another benchmark, that's the one we measure ourselves by. We didn't write the specification. There's a consortium of software vendors who write that, and get together and interpret it. We just have to get in line and do our thing, and we do.

We'll be doing it again in the very near future. We'll be doing the 30 terabytes and the 100 terabytes benchmarks. And I think as it stands today, we're going to be the only company that has those benchmarks. We'll be interested to see if anybody wants to...

CI: Join the party?

AA: Join the party and go up against us.

CI: Well, I think you might get some, but we'll see.

AA: We'll see.

CI: All right. Kevin, let me bring you into the conversation. There was a good slide on where EXASOL shines, in other words, the best types of industries that you support with your technology. Go over those for me, if you would.

KC: I like to call it data-intensive and data-expensive.

CI: That's cute.

KC: It's, basically, structured big data. This is in the context of... it's relevant to the revenue of the company, directly relevant. We have companies in new media, social gaming, and social networking so their entire business model's

predicated on being able to understand the traffic of the users and convert that into advertising or subscription sales that they do.

Other industries where it's obvious are retail. Retail trades on a micro-thin margin. Especially as you move into the online where there's, all of a sudden, another 200 different variables that now impact that micro-thin margin, analytics are critical for that industry.

Financial services, I think everyone's done that for years. It's the same. There's nothing new there. And also we're starting to see a lot of traction with what we call machine-to-machine, which are like heavy industrial big data. These are machines where the cost of a part failure is astronomical. You need to be able to instantaneously analyze the impact of that and then mitigating actions against that. These are the key industries we're working in.

CI: Or, even predict when something's going on.

KC: Absolutely.

CI: I think the whole machine sensor area is right in your sweet spot. It's a lot of very complex data, but it's also a lot of very complex queries, and predictive analytics to help someone understand what is going on, right?

KC: Absolutely.

CI: All right, well Aaron, let me go back to you. You gave us a good case study. It was a fascinating one, I thought. One of your customers and I'll let you mention them by name if you choose, but one of your customers actually took Hadoop, or at least the Hive nodes, and compared the price and performance and a few other things to what you were doing with EXASOL. I'll leave it at that, and let you describe what happened.

AA: OK, I'd be happy to. The customer in question is King, who have the game Candy Crush, which I think is pretty well known right now.

CI: Yeah.

AA: They actually started with Hadoop. They've got lots of people who are used to working with open-source, a lot of programmers, a lot of data scientists.

So that's part of the way they think. But their requirements were becoming more and more onerous, and difficult to fulfill with what they had. At one point, they tried to calculate what they would have to invest in Hive and Hadoop to get to the level of performance that they need. They found they would need 420 Hive nodes to do what 4 EXASOL nodes can do. So when you do the math on that, particularly with respect to...

CI: Over a hundred x.

AA: Yeah, you know it was a no-brainer for them really. So, in the end they ended up taking us. But, again, not instead of Hadoop, on top of Hadoop. So we integrate very closely with them, and I think we'll continue to do so.

CI: That was something that you made a point of. What was the reason for you residing alongside of Hadoop? Were there particular types of problems that fit your world better than the Hadoop world? And, of course, just the opposite, are there types of data that fit Hadoop that won't work in yours?

AA: There are both. You have the types of data. So, they collect a lot of data which you don't typically fit into a relational database management system. So they have that in Hadoop obviously.

CI: That would be the more variably-structured comments.

AA: Unstructured, semi-structured data.

CI: Yeah.

AA: Obviously, they have lots of processes that they have enough time to work through Hadoop. But, in the kind of fast-moving world that we have today, and the speed that games are played at, they have to adapt to the way customers behave. You need to have a system, which is much more agile, much more flexible, and allows real time analytics and that's what we do. So, obviously, they split their workloads between Hadoop and us, depending on the requirements of what they have to do each day.

CI: That's such a good story.

AA: I like it, too.

CI: It's a nice example of when to use one versus the other.

AA: Coexistence, exactly.

CI: It's a very confusing thing for a lot of people to get their heads around.

AA: By the way, it's not as if King -- I mean, they're very savvy, they're very cutting edge -- but it's not if they're completely clear on the best strategy for it, as well. That's something we're working on with them.

CI: Yeah, excellent. Alright, Kevin, back to you. There were a couple of other case studies. If you could just give briefly give me a quick overview of some of those cases, as well.

KC: Yeah sure as I said these are companies that are dealing with huge volumes of data. They're making big bets on, what I like to call, the "public cloud" infrastructure of the future, right. You see so much of it with Google's, and but this is happening in advertising technology, network management of WiFi spots.

For example, we have a great client, myThings. They specialize in advertising technology. They deliver over five billion impressions a month to end-users who are going through their different advertising websites. And this is about 20 terabytes a day they go through. They're analyzing it to deliver this. So it's a massive production of data to deliver the ultimate revenue generating result.

Also, Wave Fi is another great customer. They are monitoring over 200 million WiFi hotspots, globally, in a database, so that they can identify what are the usage patterns, if there are any bottlenecks that the network provider should be thinking about fixing.

Or, if you're a content provider, such as Netflix, you can make decisions about additional services you may want to purchase in certain locations, in order to make sure the end-user experience is still what you want. Those are two other great examples.

CI: All right, let me stay with you, Kevin, for just a moment longer. One of the things that was fascinating is you've kind hit your stride now. Now is the time for analytics, for analytic databases. The company's growing. In fact, you

mentioned something like 20 open slots that you're looking to hire, at least, in Europe. That doesn't even include the rest of the world.

What's the go to market strategy then? With this massive growth and this push for analytics, what's EXASOL's strategy?

KC: As you mentioned before, it's an unknown gem. It's a hidden treasure. Fourteen years of development in this solution, and it just happens to be exactly the right technology for the right time today.

Our number one top priority is raising awareness for the company, first and foremost. How could anyone buy us if they don't know who we are and what we're doing? We're investing heavily in public relations and outreach to folks like the BBBT and other important analyst communities.

Next, we're looking at, "How can we expand our channels through partners?" OEM partners, as I mentioned, are a great way for us to reach even more companies. We're also, right now, specifically in the US and some of our other markets, investing in early first customers. We're excited to be here and bring this technology to these people, and we're willing to invest in those customers to make that happen.

CI: You've opened, or will be opening, some offices?

KC: Yes, absolutely. We're opening offices right now in the United States, Israel, Sao Paulo, and London.

CI: Excellent.

KC: We even have significant activities going on in Australia and Russia as well.

CI: Excellent. You're worldwide now, for sure. All right, Aaron, back to you. Let's get into a little bit of the technical wizardry, what's under the hood and so forth. The technical features, they're called EXASolution. That's your platform. So if you don't mind, just talk through some of the wizardry that's there.

AA: I'll be happy to. I think the basic fundaments are the fact that it is in-memory. That means you can fit the whole dataset you want to analyze in-memory.

You don't have to, but you can. We have a very smart way of deciding which data needs to be in-memory and which can be disk-based.

So that's something that you have to have in mind when you want to optimize your hardware investment. We're columnar-stored and we're massively parallel. Those are probably the three most important features.

CI: They certainly are the ones that everyone's thinking about.

AA: They're the ones you need to have for...

CI: Compression maybe is a fifth one there.

AA: We have a very smart compression algorithm, as well. One interesting aspect about that is the fact that the compressed data, once it's gone into memory, doesn't need to be decompressed before you can analyze it.

In each level, what we do, we have a number of very clever technologies, which we've managed to combine in a way that I think is very smart. It's all tuned and designed to work with each other in a very good way. A lot of it is fully automated, so it's a paradigm change for many companies who are used to tuning the database every day before they run a query.

What we've also done, with a view to analytics, is to start integrating lots of functionalities into the database, such as AH, Lua, Python, Skyline Analytics, Graph Analytics, which gives a great platform for data scientists to work in a way that they couldn't do before.

CI: Well, it does. It gives them a very good way to be able to ask any question. Right? There's no restriction on the types of questions.

AA: Absolutely none. No constraints.

CI: Again, what Kevin said was it's this very complex kinds of queries. It's not necessarily the size of the data, although you can handle terabytes, if not petabytes, of data. It's more the complexity of the queries against that data that make it so very difficult for other technologies to handle, right?

AA: Well very much so, and I think the two aspects of it -- being in-memory, which is where you have amazing speed advantages over disk-based solutions, and it being massively parallel, gives you amazing possibilities.

We think in parallel. We parallelize each and every call, down to every level, and all the way up into the cluster. So, by parallelizing it, we create a massive super-computer, the kind of thing that people used to dream of, except couldn't afford.

Now you can do it with simple commodity hardware. You can build a cluster of 100, 200 nodes if you want. The amount of data you can put in there... and then each and every query is fully distributed over the cluster, and because it's in RAM it comes back in microseconds. You have your answer at a speed...

CI: Immediately.

AA: ...you couldn't get in any other solution.

CI: It's fascinating. It really is fascinating technology. All right, Kevin, let's end with you. We've got about a minute left. There were a couple of really good slides on the various forms of proofs of concepts and also the various analytic forms that you offer. If you don't mind, take a bit and talk to me about POCs and the different kinds of analytics.

KC: Absolutely, you know we want to put our money where our mouth is. We do that by setting up a test system for a client to try it out and compare for themselves. It's a new technology. It's nothing that's been seen out there before. A lot of people are talking about it right now but they're years off from there, but we're there today and we want people to try it out, so we offer the opportunity. We'll send servers over to the client where they can test it out. We have a relationship with Dell where we can set up, at any Dell Solutions Center, a trial for a customer at their servers and then even use their meeting rooms.

Lastly, we have a cloud facility where you can create an account for yourself and start testing out your own scripts -- load your data, test your scripts, and see what you think. The ultimate solution is delivered as obviously a traditional software that you can deploy on your own hardware.

We also offer an appliance model where we'll manage the whole hardware-software configuration for you. We partner with Dell, in

particular, but with other hardware vendors. Last, we have our cloud solution, EXACloud, where you can do a monthly subscription for a fixed amount of data storage for analytic consumption.

CI: Alright, excellent. You've got it all covered it sounds like to me.

KC: Good. Thank you.

CI: Alright, that's it for this edition of the BBBT podcast. Again, I'm Claudia Imhoff, and it's been a great pleasure -- I really enjoyed this -- to speak with Aaron Auld and Kevin Cox at EXASOL today, so thank you both for speaking with me.

AA: Thank you very much. It's been a great experience for us, too.

KC: Awesome.

CI: I hope you enjoyed today's podcast. You'll find more podcasts from other vendors at our website, www.bbbt.us. If you want to read more about today's session please search for our hashtag on Twitter. That's #BBBT. Please join me again for another interview. Good-bye and good business.