



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: | October 23, 2015 |
| Host: | Claudia Imhoff , President, BBBT |
| Guest(s): | Aaron Auld , CEO Sean Jackson , CMO Jens Graupmann , VP of Product Management |
| Run time: | 00:13:01 |
| 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, Sean Jackson, and Jens Graupmann. Aaron is the CEO, Sean is the CMO, and Jens is the Vice-President of Product Management for EXASOL. Welcome to you all.

Sean Jackson: Hi, Claudia. It's great to be here.

Aaron Auld: Thank you. Thanks for having us.

Jens Graupmann: Thank you.

CI: Let me start with you, Sean, quickly, an overview of why you do what you do, and how do you do it at EXASOL.

SJ: EXASOL is the developer of the world's fastest in-memory analytic database. It was conceived at the very start, and built from the ground up, to handle large data volumes and analyze them, and help businesses effectively get value from their data.

We do it through a number of ways.

There are a number of features in the product that allow us to get the performance:

One of those is that the database is in-memory. We also have an MPP infrastructure, massively parallel processing, so everything is parallelized in the database. We use compression, ACE columnar storage. We also have in-database analytics, as well.



That combination of those features allows us to then go to market and say that we have the world's fastest database.

Cl: You know, everybody comes in and says they are the fastest database or the fastest something or other, you guys are actually one of the few that I believe.

SJ: Thank you for that.

Cl: I'll tell you why. It's because of the TPC-H benchmarks. For years now, what, seven years, you've pretty much been head and shoulders above anyone else that tries the TPC-H benchmark. Is that right?

SJ: Absolutely. It's a benchmark that is an independent one. It's not something that EXASOL goes out with, and we lead the transactional processing cancels TPC-H benchmark, ranging from 100 gigabytes right up to 100 terabytes.

It's not just performance, so they don't just look at the database and say, "OK. Is it the fastest? Does it satisfy analytic queries the most rapidly?" What they look at also is price performance, through a metric they call query per hour. In that metric, we lead the charge there, as well.

This is all being independently audited, and it's all available for people to look at, at the TPC-H website, which is tpc.org, and people can go there on their leisure and look at this stuff themselves.

We're very proud of that because there really are no other independently audited benchmarks out there that tell people which is the best, and which is the fastest, which is the greatest.

Cl: Yeah, and you do top them. All right, Jens, let me bring you into the conversation. How do customers deploy your technology?

JG: We cover the whole range, actually. Our standard product is still standard software that you install on your premises, but we can also offer software preinstalled as a custom appliance, or we support virtually all of the main cloud vendors, like Amazon AWS, or Azure, or Bigstep in the UK.



CI: Aaron, let me talk to you a little bit about your company then. What do you consider the strengths of EXASOL, and where is your ideal market for this fastest database?

AA: The fact that we've been able to develop the product for 15 years now has allowed us to combine hundreds of technologies in a way that works like clockwork, and that's what gives us the impressive performance that we've managed to build.

Scalability is an important factor. You can start with one node and you can increase up to 100 or more nodes.

Agility, flexibility, the fact that it's easy to deploy, easy to integrate, but I think one of the most important ones nowadays is that it's fully automated, self-tuning, and that means that customers don't have to worry about it.

Where do these strengths take us, that's your question about ideal markets. One is definitely data driven markets, companies that rely on data to build their business models and make decisions. For them, it's a no-brainer.

But also companies that have large legacy systems and are running into problems in various areas of their organization, they can use EXASOL to solve these problems without actually ripping and replacing the existing systems.

We're also building very specific solutions to address problems that we've identified in the market, such as for Tableau Turbo, for Tableau users.

Finally, we think we can help data scientists and people who like to experiment with data with all the analytical features that we've built in.

CI: I'm going to add one more, because I think it's important, as well. The one that I would add is your transparency, about price, about who you are, about what you do, and where you do it, talk a little bit about the transparency of the company.



AA: That's a very good point. We looked at the way we position our product, and we realized people don't really understand what they're getting, and what it might cost.

So we said, "Let's just do away with any kind of questions that people have with that, and put the prices out there, so that people understand what they're getting and what they're going to have to pay for it."

Why dance around the issue if you can just lay it on the table. It's going to come down to pricing at one point.

CI: Yeah. I love that about the company.

Sean, back to you. He mentioned customers, if you don't mind, tell me a couple of stories about your customers.

SJ: Absolutely, I'd be glad to. As Aaron said, we help businesses that want to be data-driven, that want to get value from data, want to build businesses and office solutions, and proposals, and offerings to the marketplace, which are very data centric. A couple of those are as follows.

From a branding perspective, the biggest customer we have is the makers of "Candy Crush Saga," which is called King Digital Entertainment.

What they do with EXASOL is that they're looking at how people play Candy Crush Saga, and one of the other 100 and odd games that they have in their portfolio, to understand how they're playing, how they can incentive the players to spend more in the apps, how they can help people jump from level 32 to level 34, for instance.

But also, how to make the games challenging enough, but then not too hard, because you don't want to have a game that's too easy or too hard, because you just won't have anyone playing it.

They're a cache-rich company. They understand that they have a lot of data that they can analyze and mine. In that instance, what they're doing is they've got an instance of Cloudera, so they've a massive Hadoop cluster where they're doing all the data storage and data processing.



But to do the data analytics and to really analyze player behavior and the player experience, what they're doing is they're putting that data into EXASOL and analyzing it in the database there.

That's a really good case study and a great use case.

Cl: I have you to thank for my addiction to Candy Crush?

SJ: Absolutely. If you ever get stuck then just rest assured that EXASOL is there in the background and will figure it out.

Cl: It will help me out.

SJ: We'll be back helping you get to the next level.

Another good case study that we have is an online dating firm, based in London, as well, in the UK. They're also a user of Hadoop. What they're wanting to do is basically understand how people are interacting and give them a great experience using their online dating platform.

As I've said before, it's a great event for us as EXASOL to say that we help people meet people, and make of that what you will.

SJ: Another good case study is an online retailer, called Zalando. It's a lesser well-known entity perhaps, in the United States, but in Europe, they are very well known. They've got over 14 million users. I think they stopped counting, but they have about 150,000 products in their online retail shops.

What they want to do is make sure that they've got enough stock and they've got enough availability of things for their buyers to buy, and also really optimize the returns process, optimize the buying process, and just make people come back to them as opposed to their competition.

Cl: It's so important to match the supply chain with the demand chain, and that's what you're talking about.

SJ: Absolutely, and that's the beauty of EXASOL, is that it's not just for one specific use case, you can use it for anything where data is concerned, where people are trying to analyze data very quickly.



CI: A quick question, Jens, to you. It was announced that EXASOL is an in-memory database, so a lot of these companies have a lot of data. We're talking hundreds of terabytes and on up. Does all of that data have to be in-memory?

JG: In an ideal world with unlimited resources, all data would be stored in-memory, of course, but memory is still a cost factor. EXASOL uses intelligent main memory processing, so it clearly knows about hot and cold data.

By using these intelligent algorithms, as a rule of thumb, it's mostly sufficient if you just use main memory to raw data ratio of 1/10. That means you have to just have approximately 10 percent of your data in main memory, it's sufficient for most customers.

CI: You showed us an example of several customers. On average, it wasn't even 10 percent that was in memory that they were using?

JG: Yes, exactly, some of our customers even use a ratio of 1/20 up to 1/40, and they're quite happy with the performance of our database.

CI: Very impressive. I have to admit.

All right, back to you Aaron. We were just at the Tableau conference, and you were all over it. Your booth was there, a lot of excitement in your booth. You have a good relationship with Tableau. Why don't you tell me about it?

AA: It's a very good relationship, and we're very excited about it. It does a number of things for us, and obviously we do a lot for Tableau and Tableau users, in particular, but for some time we've seen analytics being decentralized, and the power moving to the users who are at the forefront of what's actually happening in business.

As that happens the requirements are growing... impatience with central IT departments taking months to deliver any kind of capabilities.

The Tableau users, they're very creative people. They create great visualizations, and it allows us to really showcase the power of the product



and the technology. If you want to work as fast as you can think, then you need that kind of performance.

The other thing that's good about the relationship is, at Tableau they're obviously growing very, very fast, but the whole culture is to enable, and they work very closely with us, which a lot of the old established vendors didn't really do, so this has given us a great amount of awareness.

I think as they advance into the market the whole London expand strategy will work for us very well, as well.

CI: Can you explain a little bit about the Tableau Turbo, that you mentioned earlier, but you didn't really talk about it much?

AA: The idea of the Tableau Turbo is a lot of the decisions about buying technology have, so far, been with the central IT department, but users want to make these decisions themselves, as well, so we had to price it in a way that would fit into their budgets and into departmental budgets.

What we did was we took our full-blown system and said, "You can have this at an extremely competitive price, as long as you're using it for Tableau." This is an important message. That's what we were doing at the conference was explaining that you're getting an amazing system, amazing performance, for an amazing price, to power your visualizations.

CI: Very exciting. Last question, to you, Sean, if I wanted to test out your product, for example, before I buy it, how would I do that?

SJ: The great thing there, it's incredibly simple to test EXASOL. We go to market with a message that we say, "It's easy to implement, easy to install, easy to get started," and we put money behind that.

Simply go to exasol.com/testdrive, register, and you can download a free copy of our software. It's the community edition. It's licensed for up to 100 gigabytes of raw data, or you can spit up an instance of EXASOL in our private cloud, called EXAcloud. Or, you could actually register for a free POC, and we'll come to you, and we'll work with you to make sure that you get the best benefits out of the technology, so it couldn't be simpler than that.



Cl: It's fully functional, right? You're not limiting the functionality?

SJ: It's a fully functional instance of the database. It's a single node only, and like I say, it's licensed up to 100 gigs. Anything else above that, then obviously we'd be more than willing to talk to you about it.

Cl: Sounds good to me.

Well we're out of time, so that's it for this edition of the BBBT podcast. Again, I'm Claudia Imhoff, and it's been a great pleasure to speak with my friends, Aaron Auld, Sean Jackson, and Jens Graupmann of EXASOL, today. Thank you all for speaking with me.

SJ: Thanks Claudia.

AA: Thanks for having us Claudia.

JG: Thank you.

Cl: 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!