



BBBT Podcast Transcript



About the BBT

The Boulder Business Intelligence Brain Trust, or BBT, 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 BBT provides a variety of services, centered around vendor presentations.

For more, see www.bbt.us

Vendor:	ThoughtSpot
Date recorded:	06-01-2016
Host:	Claudia Imhoff , Founder, BBT
Guest(s):	Anand Raghavan , Senior Director of Product Management Joe Abbott , Technical Marketing Engineer
Run time:	00:17:14



Claudia: 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 Anand Raghavan and Joe Abbott. Anand is the Senior Director of Product Management, and Joe is the Technical Marketing Engineer for ThoughtSpot. Welcome to you both.

Joe: Thank you all for having us.

Anand: Thank you for having us.

Claudia: Anand, let me start with you. Please, if you don't mind, give me an overview of ThoughtSpot, the company.

Anand: Yeah. ThoughtSpot was started about four years ago. We are based in Palo Alto, California. We have built a search-driven analytics product. The company has raised around \$90 million in funding, so far. We just announced our \$50 million Series C a couple of weeks ago.

About 100, 120 employees as of now. We've been doubling in headcount every year, and excited to be here and excited to talk about the company.

Claudia: I'm excited to hear about it. Now let's start with...you say the company supports search-driven analytics. That's your differentiator, is that search part of it. Why is this a better way to get to analytics than the traditional way?



Anand: Traditionally, if you look at our personal lives, there was a time in our personal lives where if we needed access to information, we had to go to a stockbroker, or we had to go to a travel agent, or we had to go to a librarian, but if you look at the last 20 years, companies like Google, LinkedIn, Facebook, Amazon have changed that model.

For 99 percent of our use cases in our personal lives, we get access to data through a search engine and we know how to use it. If you look at our professional lives, we are still stuck in the world of dashboards.

Dashboards serve a certain purpose, where I just need to consume information that somebody has already produced for me, but usually it's not just about the first question that you get an answer through a dashboard. It's about the second question, and the third question, and the fourth question.

Our brain thinks in questions and answers. As our personal lives have shown, search is a very natural interface for us to think about in terms of questions and answers. What we have seen in our personal lives is that if you get access to data through a search engine, not only are you getting access faster, you're going to use it more. You're going to look for more data through that interface.

Search as an interface and search-driven analytics as a mechanism to get to your data allows you to get answers quicker and enables you to actually ask more questions, because the interface is very easy and it works at the speed of your mind.

Claudia: It's interesting. I have a 23-year-old daughter. Believe me, she is very search savvy. In fact, if she ever does go to work for a corporate company, I'm sure the first question she's going to ask is, "Where's the corporate Google?" You've given her that ability, or you've given your customers that wonderful model of search, which they're already familiar with.



Anand: That's exactly what we were thinking of when we started the company, that for 20 years, people have been trained on how to use a search interface and that expectations on a search interface are also much different.

When you think of reports and think of walking up to your BI team and asking them for a report, the mental model you have is, "OK. This is going to take about a week," and I'll wait; but, if you put a search box in front of me, I want an answer in less than a second.

That's a very tall order for us to live up to in terms of expectations, but that sets up a challenge that very few companies out there can solve, so we have a huge opportunity.

Claudia: Yeah, you do. In fact, that leads into my next question, which are what are the objectives supported by ThoughtSpot for analytics?

Anand: When we started the company, there were a few goals that were clear for us from the get-go. It had to be instant answers. It had to be sub-second response times, independent of how much data is there in the system.

It had to be a very easy UI that I don't need to go sit in a three-day training class to know how to use. It had to be able to scale across all the data within the enterprise. While the interface was easy, it had to scale across all of the data within the enterprise.

It had to be a hundred percent accurate. It had to give me answers that I could trust. It had to work with complex schemas. It's not just a toy that is built with a small data set. It had to work on real data out there.

Those are all the goals that we had set for ourselves. The way it manifests itself in terms of benefits that the end user gets from us, one is the instant answers, sub-second response times, very often 100-millisecond, 200-millisecond response times, for any type of data that's out there.



The second benefit we talk about is the ability to scale to all the data within the company, to hundreds of terabytes of data and to thousands of users. It's not just a few users on the platform.

The last benefit is you want to be able to be up and running very fast. You don't want to wait for 6 months, 9 months, 12 months to have this product deployed. You want to be able to deploy quickly and get up and running. Those are the key benefits that our users get from our product.

Claudia: Very nice. All right, Joe, your turn. How does the search-driven product work?

Joe: As Anand outlined some of the key benefits we are trying to get to end users, you can tell there's a lot of sophistication going on under the hood to get to answers for business users. A business user goes to the search bar and types in regular language that they're used to using in their everyday lives. Things like sales last year, quarterly, store region, something like that, for example.

What happens is, as users type every single character, you get search suggestions that look a lot like Google's PageRank. We call it DataRank.

What happens is that we take an algorithm that combines company usage -- the usage history of you, your organization and the groups that you're a part of. We combine that with the security rules that are applied to you as a user and the groups that you are a part of. We sort those based on what would be relevant to you.

Users are getting guided suggestions as they type each character. Users select each of those words. As calculations are being done under the hood, charts appear that are actually from calculations, not prebuilt reports or dashboards.

Sales quarterly would show a line graph plotted over time. As you type something in like store region, ThoughtSpot will automatically understand that you're trying to slice your data and you want to see



something that is maybe multiple lines with different colors that show every store region.

Beyond that, what's happening is we have an analytics server, an in-memory calculation engine that is understanding what it means to calculate sales last year or quarterly, for example. Those are being translated to queries that are run on our analytics server. Those are being run on in-memory relational data cache. That's where you get the 250-millisecond response times.

Claudia: You had other features that I thought were fascinating. One of them was the ability to look at a report and ask, "Well, how was this created?" That rollback kind of idea, right? Can you explain a little about that?

Joe: Right. One of the things we're really trying to do is increase adoption within the enterprise. One feature that's more recent is our search replay feature. What that allows users to do is look at a chart that's part of a dashboard.

They may have the question, "How was this built?" They can click the search replay button and see exactly which data sources it's coming from, how the search was put together by whoever built that chart, and then how the chart was configured itself after the search was all put together.

That's one thing that, if a user who's new to the product and wants to know how to get started and build a very simple chart that they find in a dashboard that's relevant to them, they can simply click that button and see exactly how it was made.

Claudia: The other brilliance is the ability to, as you said, when you get the query ThoughtSpot will actually suggest the type of graph or visualization that perhaps this data should be displayed in, right?

Joe: Right. That is a bit about...ThoughtSpot understands what your data is made up of. It understands your data model and your schema. Based on what types of keywords you're typing in, ThoughtSpot understands what to plot where. Part of the time it takes to build



reports in other tools, you have to know what goes in on what axis, what goes on legends, how to label things correctly.

ThoughtSpot takes away all of the need to know how to do that so that business users can get quick answers from their data without having to worry about all the specifics of customizing charts and getting it exactly how they need to look. It's about speeds and sites, really.

Claudia: I think the analogy that we all came up with was you give the users guide rails. You don't let them just wander around, but they actually get guided in this creation of these reports.

Joe: Exactly. That's the second piece to it, is that we have these objects in our product called worksheets. Essentially what that is is a very thin layer on top of all of your tables that sit in the data cache.

That is something that you can do to make it easier for end users to find their answers more quickly. You can select a set of columns for users to know exactly kind of what set of questions they want to ask. You can create synonyms.

You can do some basic modeling -- say, how certain columns should be treated when graphed -- and all those things that are done by an analyst or somebody that is a power user makes it very easy for a regular end user to type in a search and get a chart that they find effective.

Claudia: Now, this doesn't come with just a thin interface into the data, I suspect. If you don't mind, spend a little bit of time talking about the architecture under the covers.

Joe: Sitting up at the very top is our relational search engine. That's where a lot of the suggestions and the usage-based ranking or the search history is stored. This comes from our search index. That search index is what allows users to find their answers pretty quickly and get evolving suggestions as they type.



Underneath that, you have a visualization and a collaboration layer. That's what allows automatic charts to be produced. That's the technology that understands what makes up each of the individual search terms and knows how to plot them on the canvas.

Then you have a collaboration piece, the ability to share with users and groups and have the security rules persist as you share. Beneath that, you have an analytics server. That's what's translating these regular search queries into sequel queries that are being run on our in-memory relational cache.

That sits on top of a distributed cluster manager. At the very bottom, we have the ability to scale out to many terabytes, several terabytes. That's the full stack. Alongside that, we have connectors for the most popular data sources -- Salesforce, Oracle, SQL Server, MySQL.

On the other side, you have enterprise-class securities with things like row and column level security, object level security. Those things bubble up to the very top of the stacks. If users are not allowed to see a certain set of information as they search, they won't see that information. You can be sure that, with this platform, you're insuring security from the top down.

Claudia: Let's talk about that in a little more detail. You've already talked about the security, but you also have audit and governance capabilities, right?

Joe: That's right, yeah. Administrators have a special extra section that they can go to, to look at what are the most commonly used reports, how often users are searching, what adoption looks like within different organizations.

That's what measuring success and measuring adoption capabilities look like in ThoughtSpot, but as far as governance, we also have very fine-grained privileges that allow users to only perform a certain set of operations.



You can create a group of users that can only view data or a group of users that can edit data but not download it, and lot of different combinations of things that make it very easy to keep all of your analytics centralized within ThoughtSpot.

Then beyond that, we also have some auditing capabilities. You can see what actions are performed at an administrative level. You can track what new data enters the system. There's also some other features like stickers, and labeling, and tagging that make it easier to stamp which data sources are actually blessed data sources.

Anand: Some of the other cool things we have in our platform is, if you have LDAP, you have SAML, you have integration into all of that. You also have a very easy way for users to use our UI to bring data into our product. It's a very, very light, easy way through our UI to get data in.

The other important thing when you use a product that has a search bar is to make sure that the users can trust what they're seeing and there's transparency in terms of what is going into computing those results.

We spent a lot of time paying attention to that. We have what we call Search Inspector, which allows you to...you click on that, and it will tell you in plain English how that query was computed, how the result was computed, which tables were joined, how the filters were applied, what dates were chosen, and what led to computing that result.

This is an easy way for a business user to verify that the result was computed the right way. If they had any questions, they could show it to their BI team, who could then write SQL on another system and verify if it looks different or not.

Another important part of that is the ability to track the lineage of any of the data you're looking at. What table did it come from? Which column? When was it last updated? We transparently allow access to all of that for the user to look at and verify.



To your point about guide rails, these are the guide rails that we have in place to make sure that adoption doesn't come at the risk of loss of trust in information. Those are all features we have added to the product right from the get-go.

Claudia: All right. I'm going to stay with you for a second, Anand, for the last question. Why don't you tell me a little bit about your customers and maybe their usage of ThoughtSpot? How are they using it?

Anand: Yeah, absolutely. There's a lot of Fortune 500 customers that use us. There's plenty of mid-market companies that use us as well. Some of the typical use cases we have seen, there's a very large retailer, one of the top retailers in the world, is a customer of ours. They have all their data sitting in Teradata.

They have a traditional desktop tool that builds dashboards for them. This is their e-commerce platform, where they look at weekly product sales for the millions of SKUs that they have that they sell online.

The difficulty that they were seeing was that these are recent grads, who are merchants, who are responsible for thousands of SKUs. Each person is responsible for thousands of SKUs. There's only so many dashboards they can look at in a week, let alone the time it takes to build those dashboards.

What they wanted ThoughtSpot to do was to provide the ability to search quickly and find sales information for all of these different SKUs for each of those merchants. They have started with a few 10 people about a year ago to 600, 700 people using us on a weekly basis, writing about 8,000 queries a week.

It's been a really good success story for us. They have bought multiple times from us as their needs grew and as they had more data to put on ThoughtSpot.

There's another customer, which is a very large wireless mobile provider. They use us to look at churn data to see how many customers have churned because of a campaign that they ran.



Believe it or not, before they had ThoughtSpot, it was a three-month process for them to get the data that they could look at to see how a campaign that they ran three months ago performed from a churn perspective.

In less than a year, they have gone from it taking three months to daily churn analysis now, to the point where the users are asking for hourly churn analysis, the same users who were OK with a three-month latency before.

That's what happens when you put a search engine in front of people. Your data consumption patterns change and your expectations of your data change. That's very encouraging for us because that's what we want people to do. We want people to make quick decisions off of their data. Those are a couple of examples. I can go on, but I'm sure we have time limitations.

Claudia: We are, unfortunately, almost out of time, but I think it's fascinating, what you just said, that their patterns of actually querying the data change based on the ease and the comfort that they have with your technology.

Anand: Right. It's been very, very heartening to see usage increase across all four customers, where it starts with the BI team and probably a few people in the line of business. They realize that this is actually really interesting and it's useful.

You see week over week, you see a lot of adoption at these customers. It helps that we are priced for adoption. We don't charge per user licenses. You buy a certain amount of data volume from us. You can add as many user logins as you want. Our incentives are aligned with our customers. You buy more from us only if there's more adoption at your company site. We are built for adoption.

Claudia: Yeah, that's great. Unfortunately, we are out of time now but thank you so much. That's it for this edition of the BBT Podcast. Again, I'm



Claudia Imhoff. It's been such a pleasure to speak with Anand Raghavan and Joe Abbott of ThoughtSpot today. Thank you both.

Joe: Thank you.

Anand: Thank you.

Claudia: 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!