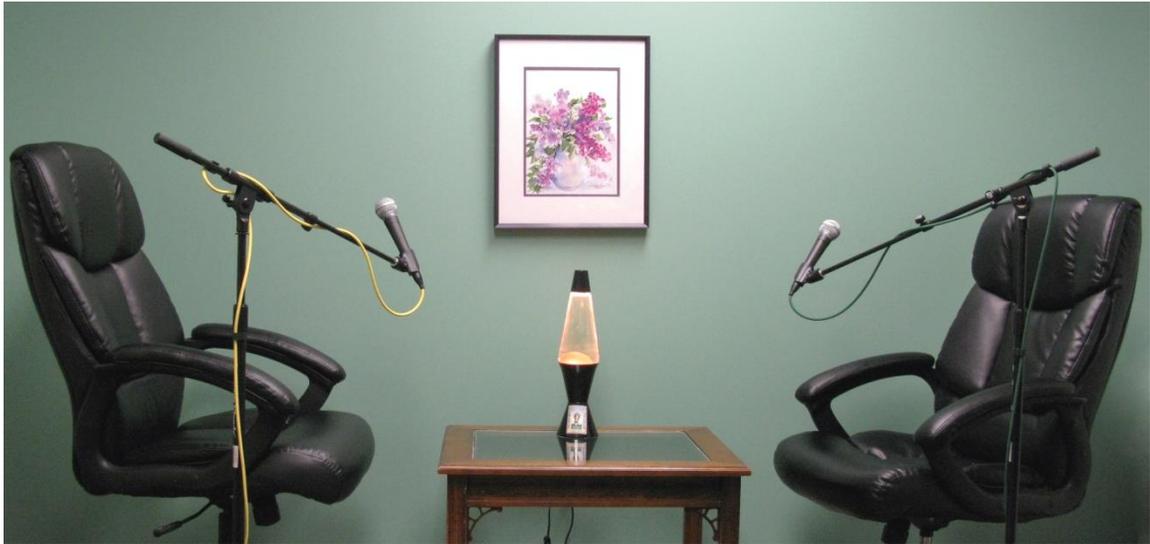




## BBBT Podcast Transcript



### 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](http://www.bbbt.us).

<b>Vendor:</b>	<b>Teradata</b>
<b>Date recorded:</b>	<b>November 8, 2013</b>
<b>Host:</b>	<b>Claudia Imhoff</b> , President, BBBT
<b>Guest(s):</b>	<b>David Schrader</b> , Vice President, Products and Services Marketing
	<b>Chris Twogood</b> , Director of Marketing, Unified Data Architecture
<b>Run time:</b>	<b>00:16:52</b>
<b>Audio link:</b>	<a href="#">Podcast</a>
<b>Transcript:</b>	[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 so pleased to introduce my guests today. They are my good friends Chris Twogood and Dave Schrader. Chris is the Vice President of Teradata Products and Services Marketing, and Dave is the Director of the Unified Data Architecture Marketing for Teradata. So, welcome to you both.

Chris Twogood: It's great to be here, thanks for having us.

David Schrader: I always love coming back to Colorado, thanks for inviting us.

CI: I love having you guys here. Let me start with you Dave. You started off with some key messages about the Unified Data Architecture. First of all, let's at least give me a high level. What is the Unified Data Architecture, or the UDA, and what are these key messages?

DS: The Unified Data Architecture is Teradata's response to the big data phenomenon. It can be horribly complicated. I give a lot of talks at conferences where people's eyes are glazing over because there are so many piece parts, so many data types, and so many things that you can do.

Our main corporate goal is to handle any kind of data, no matter what the type is, provide tools and analytics for transforming data into insights and then integrating those all back so you can take some action.

We view it as an interesting phenomenon, a little overhyped, at the moment. But the key point is that it's still just an evolution not a revolution. I've been around a long time. I've seen the days when object oriented databases were going to replace and solve world hunger.

CI: I think we've all been down those roads.



DS: When it gets overhyped, it's just too much. We try to take a very practical, down to earth, well-engineered approach to integrating all the components, the connectors that you need, so that you can build a system. Then add in any kind of data, any kind of analytic, and then take action so you can get some business value out of the new technologies.

CI: Well, there's also some papers, that you've got, some assets that people can look into if they want to read more.

DS: We've written a variety of interesting things. There's a new paper called, "Big Data, Unified Data, Architecture in Action," that got posted last week.

For business people, a nice, Scott Gnau, who's the President of TeraData Labs, paper to be used to educate executives. For people who want to get down in the details, we have plenty of things that we produced or analysts have written, with our help.

If people go to [teradata.com](http://teradata.com), there's a lot of information there that they can use to educate themselves.

CI: All right, well, let's talk a little bit about your customers. You mentioned a lot of interesting ways that customers are using the UDA. What are some of these?

DS: I get the opportunity to listen to a lot of our internal readouts of what people are doing in the proof of concepts. I find them to fit into a variety of buckets.

Some people are trying to analyze customer behavior. In addition to the traditional demographic styles. You can look at the behavior with big data. It could be any kind of behavior. It might be a customer who got a bill, and then they went on the website, and they took a look at the details. Maybe it's a fee from a bank. They spent two minutes on the website in the fee section, then they switched and called the care center.

Now they are call center notes, things that maybe traditionally you didn't put in the data warehouse. You can start pulling them together. You can start doing voice to text and word spotting. All of those new evidence points. The amount of time on the fee page, the tone, even, of the call into



customer care. They can be scored and added to the traditional predictive models to make your customer treatments then, much more accurate, more nuanced, precise, and personalized.

In customer management, that together with a hot, hot, hot area of sentiment analytics. What are people Tweeting? Are they talking about my company? What are they talking about? Is it good or negative sentiments? Who are the influencers? All those kind of things get lumped into customer management, big data.

The second big category would be, not just customers, but things. The Internet of things. There's a report by Cisco that I read, on the airplane coming here, that said that, by the year 2015, there will be 15 billion sensors in the world, kicking out signals.

Certainly, everything from automobiles, airplanes, have all kinds of sensors. When you analyze the sequence of behaviors and reports that are coming out of those devices, you can do things like spot problems. A machine might be breaking down or having a major problem. It might be a minor problem. You can do better prediction on the parts and the replacements and what you need to do in warranty situations.

A huge number of things. We have the benefit of being a big worldwide company of seeing of a big panorama of the uses of big data.

CI: Very quickly, if you don't mind, you mentioned four steps to making all of this actionable. Tell me about the four steps.

DS: The first step is you have to capture all the data. There may be a multiple of systems. This is like EDW Enterprise Data Warehouse 101. You've got to get the data into place so you can analyze it.

The unified data architecture, we have this new area called the Discovery Platform. Instead of doing it on the data warehouse, because the data might be dirty, and you might be doing quick and dirty experiments, you don't need the high degree of precision. So step one is pull the data into one place.



Step two is then to pull the information from the different sources into order by time. That way you can see the time sequence of what happened in what order, and put it back in order for selected customers or machines of interest. Step three would then be to visualize the patterns that you're seeing.

Once you've discovered interesting patterns and insights then you need to fold it back into the mainline systems and operationalize the insights. That pattern holds over and over. It doesn't really matter if it's customer or Internet of Things, types of equipment, cases.

CI: Really interesting, because it brings in the ability to take those models, those predictions, whatever they are, and actually be able to act on them. That's the whole key to it, isn't it?

DS: Yeah, it's data to insights, to action, to value, and we see that pattern over, and over, and over.

CI: All right Chris, let me get some details from you. What I'd like to understand a little bit more, a bit more detail, about the unified data architecture. Maybe you could walk us through the big components there, the big parts.

CT: Yeah, I'd be happy to. It's interesting how Dave said, it's all about data insight, action, and value. I don't think that's changed. When you're doing analytics you want to be able to uncover stuff, you want to be able to discover and then put it into action, so you're driving that value for your company. I think what's happened though, with this just onslaught of all of this data, it's changed the types of architectures that we need to think about.

A lot of times this onslaught of data, there's a lot of noise in the data. What you're looking for, is you're trying to uncover the signals, or as we sometimes call the business value density. How much density is there in all that data, which we can extract from? What this drives towards is the need for a new architecture. This is our unified data architecture. It really has four key components to it.

The first component is what we call a data platform. The characteristics of the data platform is where you would capture and store all the raw data.



You would refine the data and prepare the data for doing downstream analytics. You could also use it for an analytic archive, so it's available, and the source is stored there for a long period of time, but the data is still in its rawest form.

Then what you need, is you need technology where, what I call, mere mortals can look at data, you know, without having to write Java or Python, or Perl, or Ruby...

CI: Oh, you're speaking my language now.

CT: Slow down, right? You know, where they can literally be able to invoke functions to look for things in the data. Let the data tell you the questions you're not asking. What are the patterns of data, what are the affinities, what are the things in terms of the time series? So we believe a key component of the unified data architecture is the second element, which is a discovery platform.

This is where people that understand SQL, and we call SQL the language of business, really can go through and invoke different types of technologies, different patterns, SQL, and MapReduce, and graph technologies, to just uncover data. Once they've discovered that, they then want to take and they want to enrich the data warehouse. The data warehouse is really where you bring together all the cross functional data, where you share it with the broader set of users and suppliers, and you really help operationalize the technology.

Those are three of the four, but the fourth piece is really being able to make the access to the data easy within the environment. Make it easy to manage overall architecture, and make the data movement between the different technologies very seamless.

At Teradata, we spent a lot of time adding value in each of those different four categories, so it's really simple for our customers to extract that value out of big data.

CI: Well let's drill into that just a little bit more. You're right, it's been an incredibly innovative time, and that's usually a very disruptive time. Each of these areas has their particular technological bent, if you will. For example, the



data platform, there are certain technologies that lend themselves to that environment. You mentioned Aster for the discovery one, and Teradata, I'm assuming, for the data warehouse one.

If we look at these different environments, and these disruptive technologies to, as you say, "the mere mortal," it becomes very complex. It becomes very confusing in many cases. How do you simplify this environment? You started to talk about it. How do you simplify the administration, the access to the different components and, of course, the different computing platforms themselves?

CT: Yes, it's a great question, Claudia, because if you think about it, the data platform, a great technology for that space is Hadoop, right? Certainly, Hadoop has some great characteristics for low cost, fast ingestion, ability to do very big batch processing jobs. But on the other hand, there's a lot of industry feedback about how mature is it? Is it ready for the enterprise? Is it secured? We've all heard this.

At Teradata and with our partnership with Hortonworks, we do a lot of things to put enterprise capabilities around that technology. One of the things that we do is a technology we have called SQL-H, which literally makes it transparent for the user to access data. They could be looking at tables and it represents information that's sitting in Hadoop but, for them, they don't care. They don't know where it's sitting, and with SQL-H they can just do the query.

What the technology will do underneath the covers is it'll go read the metadata layer, it'll transact it and it'll bring it forward and, on the fly, it'll process it. From a user perspective, there's some data, I want to join it with this data, and I want it to be secure, so you go through Teradata or the Aster environment to get there.

There's other things that we're doing to provide what we call a single plane of glass in which you can manage your entire system. So you can look at what is the health of the systems? How are they running? Are they running efficiently? To enable administration individuals to do that, without having to look at five different things that have different functions and usability and work in a very different kind of way. There's other things that we're doing



just to make it very proactive, and making sure that the platforms themselves are always up and running.

As a key mantra for Teradata, and it's always been our history, is we want to simplify the technology for the users and that continues with the unified data architecture.

CI: A lot of innovations. Real quickly, it sounds like, does all of this have to be on premises or is it possible that we can bring in the Cloud, for example, for a part of this or some of the unified data architecture?

CT: It's a great question because, especially with big data, some people are making strides with big data, others are still like, "OK, what does this mean for me?"

CI: In their infancy.

CT: Yes, they're in their infancy. So with Teradata, and we just announced this a few weeks ago, we introduced the Teradata Cloud.

The Teradata Cloud really has three key components to it. It has data warehouses of services. If you want to do data warehousing in the Cloud, you can do that with a full complement of services that we offer.

You can also do discovery as a service with the Aster discovery platform and using all the graph, and the MapReduce, and SQL, in the Cloud. As well as data management as a service, which is Hadoop. Now the beauty of this is it gives our customers flexibility. If they want to do a hybrid Cloud, if they want to say, "Hey, I'm going to do my big data in the Cloud, test it out before I go deploy it." or if they want to do their full production, enterprise data warehouse in the Cloud, we can do that.

I think it gives our customers a lot of flexibility around... They can do it all in Cloud. They can do it all on prem. They could do some on prem, they could do some in the Cloud. Depending on what their evolution is within their deployment model around the UDA.

CI: It's a brilliant model, it really is. OK, last question. What's coming down the pipe? What do you see in terms of the next six months? Two years? What new innovations?



CT: It's another great question because this environment is just innovating so fast.

CI: It's moving.

CT: We just saw Hadoop 2.0 introduced, and everything they're doing with Yarn, and really impacting some key elements around being able to better manage workload management. What we see as really the next level of innovation is how do you continually increase the interoperability of these different technologies. Because what you don't want to do is you don't want to say, "Oh, I want Hadoop to become a data warehouse," or "I want a data warehouse to become Hadoop," because they're really designed for key functions within each of those spaces. As you change them, you lose some of the key design principles.

What you'll be seeing from Teradata is really what we talk about around fabric based computing. This is about ensuring you have a very high speed interconnect, and then you start to look at opportunities, not just to move data, but also to push down processing for the right foundation and the right technology. Use the right tool for the right job, and leverage the engine's capabilities but, again, make it very transparent for the user.

CI: Oh that's wonderful. Not everything is a hammer, right?

CT: Exactly.

CI: OK. Well, 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 Chris Twogood and Dave Schrader of Teradata today. So, thank you both for speaking with me.

CT: Thanks again. Thanks for your time.

DS: Great to be here and see you again. Thanks.

CI: I hope you enjoyed today's podcast. You'll find more podcasts from other vendors at our web site. That's [www.boulderbibraintrust.org](http://www.boulderbibraintrust.org). If you want to learn 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!