BBBT Podcast Transcript

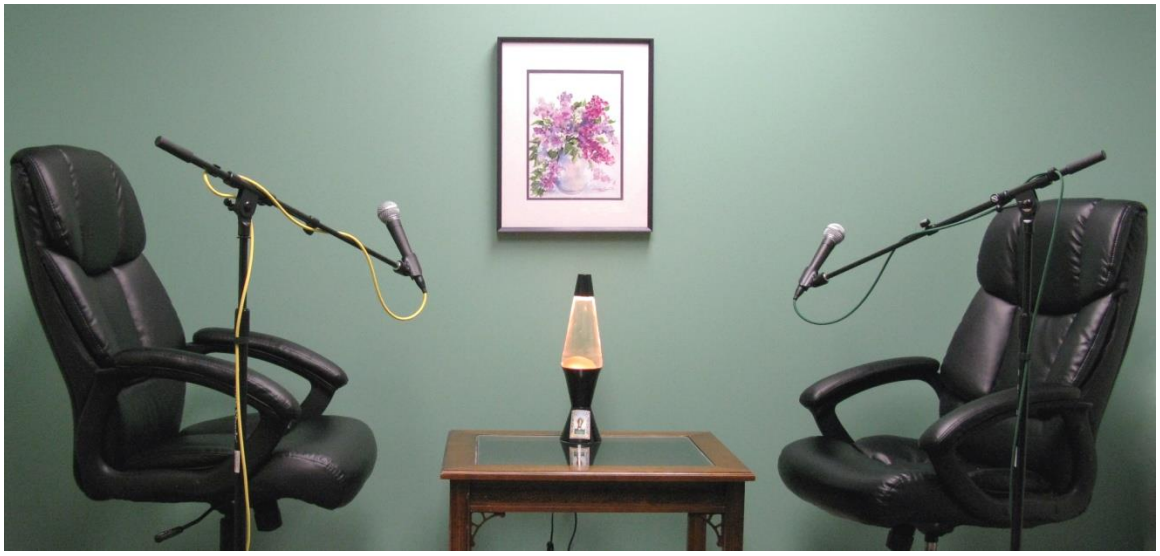

About the BBBT

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:	**Cisco Composite**
Date recorded:	**August 1, 2014**
Host:	**Claudia Imhoff**, President, BBBT
Guest(s):	**Mike Flannagan**, General Manager, Data & Analytics Business Group
	Jim Green, CTO, Data & Analytics Business Group
Run time:	**00:20:06**
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 Mike Flannagan and Jim Green.

Mike is the General Manager and Jim is the Chief Technology Officer for the Data and Analytics Business Group for Cisco. Welcome to you both.

Mike Flannagan: Thank you.

Jim Green: Thank you.

CI: Alright, Mike let's start with you. You began by describing the Cisco difference, and within the data and analytics business group at Cisco you mentioned that there are four areas of focus, if you will, four areas of interest. What are these?

MF: They're all tightly coupled, Claudia. If you think about big data and analytics, on large data sets, over many years that help you drive insights about your business and help you determine which policies you should set within your company, you then need to find a way to enforce those policies all the way out at the edge of the network, given the fact that much of our data is becoming very distributed and decision making can't always be sent all the way back into the cloud or all the way back into the data center.

Three of the four areas are really focused on those core assets, that do big statistical analysis of very large data sets and help derive business insights. The next layer down from that would be part of our prime analytics suite that's really focused on using those data and the insights from the data to instantiate policies.

Then, finally, out at the network edge, Cisco's Data in Motion products, are about embedding intelligence into edge network devices to do real time enforcement of policies, all the way out at the network edge so you can enforce those policies in your real time.

CI: Just for people who may not know what "edge" computing is, a brief description please.

MF: Taking existing network infrastructure in the form of routers and switches and embedding the ability to do compute tasks, all the way out at the branch office of a bank, all the way out at a retail store, all the way out in an oil rig, in the middle of the ocean. Being able to do computing locally there, is edge computing.

CI: It really is where the network ends and the device hooks on?

MF: Correct.

CI: Mike, continuing with you, let's talk about the Cisco building blocks or Legos, as you like to call them. It was quite a comprehensive set. I'm going to let you talk about the three layers, if you will, to this architecture.

MF: Certainly, one of those building blocks are the fourth area from the previous question, and that is data virtualization and pulling together all of the data from across the enterprise.

In a world where data is becoming increasingly distributed, data virtualization is becoming increasingly relevant. It's just not nice to have technology, but really a must-have technology, as part of everybody's information architecture.

Along with data virtualization, PRIME Analytics, Cisco Data in Motion, there is a variety of other building blocks, and you could think of those as raw ingredients. Those raw ingredients can be put together to solve a business problem, either by a very savvy end-user customer, or perhaps by a systems integration partner of Cisco's. That is the most fundamental layer where you're working with raw ingredients.

The opposite end to that spectrum, then, would be fully finished products, where Cisco has done all of the integration of a variety of different products

together with some services to create a solution to a customer's business problem.

Then in the middle, there are pre-integrated packs of products or services that begin to solve a problem, but require someone to do a little bit more work to fully solve the business problem.

If you want think of that in terms of an analogy, the one that I used earlier today in our session was baking a cake. As a fully finished product, you can walk into a bakery, pick out the cake that you'd like to eat and take it home. Other than paying for it, there's nothing that's really required for you in the work effort, and there's really no knowledge of the underlying technology, the recipe, or the cooking that you need to have.

One step down from that, then, would be the Betty Crocker cake mix box, where you understand how to bake the cake, but you don't want to have to do the shopping for the raw ingredients. You don't want to think about exactly how much flour do I need to put in, so you wanted to be a little bit easier for you but you want to do the fun part of putting it all together to make the cake.

Then, at the most fundamental level, you have grandma's recipe, you go to the grocery store with a shopping list and pick out your eggs and flour and butter and sugar. You really not only understand how to bake the cake, but also how to put the right ingredients in so that you bake the right cake at the end.

CI: Jim, let me turn to you now or bring you into the conversation. We had a very thoughtful, thought provoking session with you on the Internet of things or the Internet of Everything, as Cisco likes to put it. If you don't mind, just give our audience a brief overview of some of the things that you see coming in the Internet of Everything.

JG: There's been a lot of discussion on the Internet of things and projections about it being a remarkably big market opportunity, measured in trillions of dollars. A lot of my presentation could be described is demystifying the hype. The Internet of Things is a combination of applications, data, networking, and devices.

If you think about it, all four of those things actually exist. There's not a new major element to this. The newness is how you would combine them into a system, and how you would do so in such an extent that the programming model is easy, that the system is scalable, and it's maintainable, and lifecycle cost is controllable, and the system runs predictably.

It's more of a methodology on how you combine these things rather than a new thing itself.

CI: Which I greatly appreciated. The demystification was really so needed. It's like big data, there's so much hype around the Internet of Everything, the Internet of Things. There's so much hype around big data.

It was refreshing to see something that made sense, the different ways that we interact with things, where transitions come from, what we're going to do with all of these devices that are spewing information at us. A little scary at the same time. It also means that there's an enormous play for data virtualization in this environment, is there not?

JG: Yes. One of the things that we face is that the devices are, in many cases, geographically separated. Take a city infrastructure management, or, even worse, take a smart meter utility grid scenario where you've got the devices over a wide geographical area.

In some cases, we just don't have a lot of network bandwidth all the way out to the edge. Another example is offshore drilling for oil, where you have a lot of equipment out there. You've got a lot of money at stake. We don't have a multi-gigabit network out at the edge.

Therefore, what will be required is to capture the information and to actually look at, analyze, filter, standardize or normalize the data, then very selectively use that limited network bandwidth to pull the most-needed information back to the center so that a geologist sitting in Houston can analyze what happened in drilling activities out there in some remote part of the world today.

CI: The other thing that you said that was so refreshing as well is that the IoE, or the Internet of Everything problem is a data problem. I love that. It goes back to your statement of the technologies aren't new. We have

everything that we need. What we need to solve is the data problem. Explain what you mean by that.

JG: I'll explain it as long as you promise not to broadcast this all over Cisco, because most people at Cisco think it's a networking problem.

CI: Absolutely. I promise I won't say a word.

JG: The more you think about it, and the more practical you attempt to be, then the more certain things become obvious. Go back to the utility system. You read a meter. It gives you the load. It gives you utilization in terms of kilowatts since the last time you read it.

It gives you a timestamp, and it gives you a unique meter ID. That, maybe, is 100 bytes of data. Utility companies want to read the meters every 15 minutes. That means that you've got 96 reads per meter per day, times 100 bytes. Then you've got five million homes.

CI: I was going to say, let's start doing the math.

JG: That's right. You can actually do the math, and what you realize is that you're creating billions of bytes of data per day, which is not that big a deal to process it, but you start saying, "How many days' worth of data can I store"? What you realize is that the cost is in the storage. You really need to think about the quality of data.

"Do I store the aggregates? Do I store the raw data? Can I do the analysis at the edge, or at least as soon as possible, and then store the results of the analytics instead of the data itself"? You just are required to get practical to keep the costs from spiraling out of control. The other factor is, as we all know, the bigger the data base, the slower the query.

Is it possible to actually sort out at the edge what is really important and therefore make the data storage smaller, which then enhances your ability of looking at the data in multiple ways in a reasonable amount of time, because inquiries will actually run faster?

CI: Let's get practical here with a couple of examples. You gave us a number of examples of companies that are taking this to heart, using this model that you came up with and actually doing some really cool things.

You've got some solutions, these pre-done cakes as you would put it, Mike, that are doing some really phenomenal stuff. Give our audience just a few examples of these.

JG: One of the things I think the audience should start making a distinction is across different verticals.

For example, there is going to be a lot of discussion about the smart home. Google is in there with a big acquisition. Apple Home is in there. Microsoft will respond. Honeywell, who builds thermostats, has got to have some sort of reaction to all of this. There's going to be a lot of discussion and a lot of hype. It's something that the consumers will relate to and so there will be a lot of press around.

That's not really where the action is, because this is high-volume, low-cost, low-margin activity. It's going to be very difficult to make money on that.

In the Google world, this is eyeballs and extending Google into the home as an extension of the franchise they have on the Internet.

The other thing is the technology here is peer-to-peer, low-volume and really low computation. The examples are I've got something on my wrist, so that when I get out of bed in the morning, the shower turns on. OK, whatever. The real opportunity here...

CI: Coffee pot starts.

JG: That's right. I want it in my home, but that's not where the business is for all of us. We're going to see a lot of work in oil and gas. We're going to see a lot of work in industrial automation in the plants. We're going to see a lot of utility work in smart grids.

In this market, it's going to be interesting. The public sector is going to come in much earlier than expected with smart cities. Because even though they're constrained in terms of their budgets, resolving traffic congestion and optimizing trash collection is a huge savings in cost.

Emergency response times are a huge increase in service. If somebody gets in a car wreck, they want to know where all the ambulances are and send the nearest ambulance.

When they pick that person up, they want to determine which hospital are you going to and turn all the stoplights green.

There's going to be a bunch of activities like this that are really important and meaningful in terms of quality of life and in terms of profit lost to large corporations.

CI: Fascinating. It's just an incredible future that we're all facing, which I find very exciting. Let's get back to Cisco data virtualization.

You mentioned that there are five, I think, new strategies for the Cisco data virtualization future. What are these and why are they important to the future of the data virtualization area?

JG: The data virtualization category needs to be agile and change with the emerging times. For example, the big installations of data virtualization were originally in the financial sector. It wasn't even the retail banks.

It was the trading banks. We build a very high scalable, ultra-high-performance product that was a very technical product, because they have high intelligence and IT. They were a very demanding customer.

That served us very well as we expanded into pharmaceutical and into oil and gas, but today we see retail, manufacturing and even health care. We see companies of small- and medium-sized enterprises getting into the business of data virtualization, and they don't have the IT budgets that the investment banks did.

We need to make the product easier to use. We need to make it easier to install. We need to shorten the time between the acquisition, the training, and getting into production quickly.

We're spending a lot of effort in terms of taking all the capability of the product and rendering it in such a way that people can be more productive with it quicker. It's a lot of work on tools, data virtualization.

While we were doing this, big data grew up. We now have to all be very serious about integrating big data with the rest of our enterprise scenario.

I am, and we are, of the school that big data has a place, a very important role, but it's not a revolution. We're still going to have our OLTP data bases, which follow all the ACID properties and do queries in a very responsive way.

As we do that, we need to figure out is, what are the ways in which big data intersects IT and how do we incorporate that in such a way that people can build hybrid systems, we can build logical data warehouses with multiple physical repositories, and we can do so in a very straightforward fashion with the ease-of-use tools.

The other thing that we are seeing has changed over time is the increase in analytics. When we started getting into this business, we would inquire the VI companies about analytics. They would wave them off as the weird guys in the corner that were the minority user, but reporting and operational issues in dashboards was a real source of revenue for them. Boy, how things have changed. So as we look...

CI: For the better, I think.

JG: Yes, me too. I agree wholeheartedly. But as we look at this, we see the data scientists and data analysts being completely self-sufficient, trying to get the answers to data questions without going through a standard IT-SDLC process. We see them doing things differently, like building sandboxes. We see them actually trying to do multiple turns of looking at the data in different ways within a single 24-hour period.

That requires custom tooling for those folks, in order to support their type of behavior, their scope of work, and the way they actually approach things.

We're just trying to adjust the product, which is a very capable product, to meet the changing needs of the industry, as the industry itself evolves into new ways of working.

CI: Yeah, and as your customers themselves evolve, like you said, they're no longer doing just reporting, or just multi-dimensional analysis. They are actually doing some advanced analytics and demanding from you the ability to get at that data to do that.

JG: Yeah, I was talking to a really smart guy and one of our partners the other day. In mid-sentence, he stopped and he says, "The more I work on this, the less I'm certain as to whether or not, I'm working on big data, on analytics, or on IOT."

CI: It's all going to blur.

JG: It's all going to blur. You cannot do IOT and get the business benefits unless you're dealing with a lot of big data and you're doing analytics on all the information that's being generated by the devices.

CI: At some point, it's just going to be data.

JG: It's just going to be data.

CI: With that, we'll have to say that we're done. 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 Mike Flannagan and Jim Green of the Cisco Data and Analytics Business Group. Thanks to both of you.

JG: Thank you.

MF: Thanks, Claudia. It's always great to be here.

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!