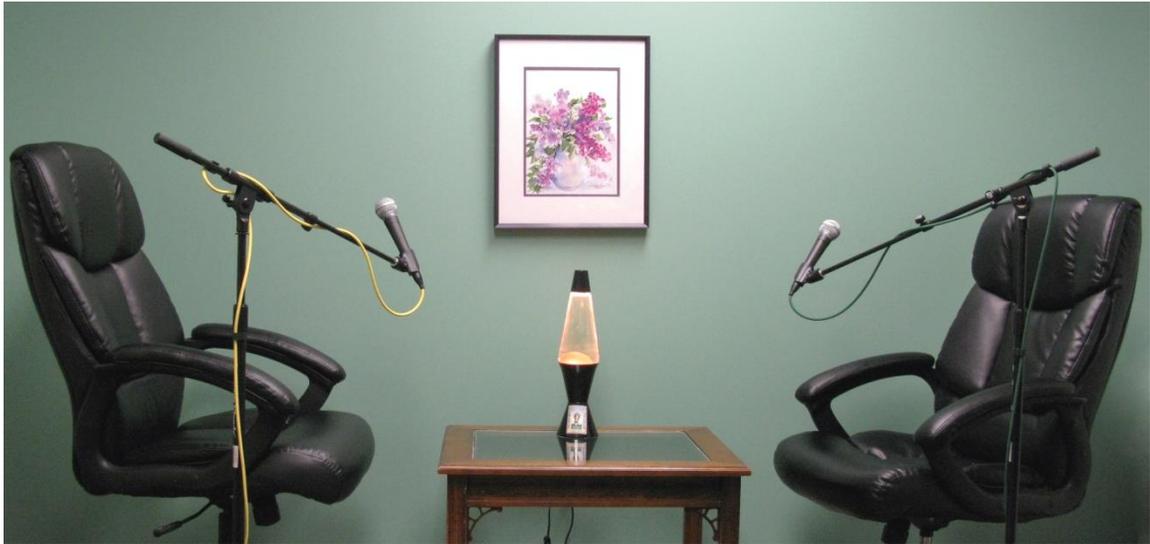




BBBT Podcast Transcript



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:	Splunk
Date recorded:	August 22, 2014
Host:	Claudia Imhoff , President, BBBT
Guest(s):	Manish Jiandani , Director of Solution Marketing
Run time:	00:17:25
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 please to introduce my guest today. He is Manish Jiandani. Manish is the Director of Solution Marketing for Business Analytics for Splunk. Welcome, Manish.

Manish Jiandani: Thank you, Claudia, it's a pleasure to be here.

CI: I'm thrilled to have you here. It's the first time Splunk has come to a Boulder BI Brain Trust. Let's start out with who is Splunk, I guess. Splunk is not necessarily a new company. In fact, you've been around for more than a decade. I suspect there are a number of people who may not know a lot about it. Why don't we start with a short overview?

MJ: Sure. So, Splunk, as a company, was formed in 2004. I think that's when the co-founders, Erik Swan and Rob Das, came together to venture on this journey of trying to help customers gain insights from their log data.

There's an interesting story about how the name came around. I think when they were validating their ideas with prospects, the feedback that they got, very often, was that looking for insights from log data was just like going down a cave wearing head lamps, and that was the genesis of the name "Splunk." Splunking is an activity you can do when you wear head lamps and you discover caves.

CI: Spelunking.

MJ: Spelunking was the term, Splunk was born out of that. Since 2004, the company has grown a lot. 2006 is when we launched our first software product.



Since then, we have around 1,200 plus employees now. It's a big company. We keep on adding new and new employees every day. The company is headquartered in San Francisco. We also have branches all across the globe. It's a publicly traded company, as well.

We have customers in 90 plus countries now, about 7,400 plus customers and that number keeps on growing, as well. We, in terms of the customer profile, we are in about 65 of Fortune 100 companies or two-thirds of the Fortune 100 companies use us.

We also have presence in over 90 plus countries today. Our largest customers are indexing about 200 terabytes of data a day. That's awesome and amazing. In terms of the business model, our biggest deployment, right now, is on premise. We have customers who are using Splunk heavily on premise, but we also have a cloud solution, and also a SAS solution.

CI: All right, that's quite a bit. Let's talk about the sweet spot for Splunk. It's what you call, "machine data." We had a big discussion about that. Why don't we start with what exactly is machine data?

MJ: Machine data is anything that is coming out of machines. This could be logs from your Web servers, logs from your applications servers.

This could be data coming in from devices in your infrastructure. This could be VMware machines, it could be routers, switches. These could be devices like your iOS devices, Android devices, variables, sensors, geolocation data.

All this data, for us, is machine data. What is unique about this data is essentially this data is always time stamped, so it's time series data. What we have found is that customers can gain a lot of valuable insights and intelligence from the machine data. If you look at the growth of the segment, machine data is probably the fastest growing, most complex and most valuable segment of big data today.

CI: Yeah, boy, it's just remarkable. It's pretty much almost all data, in many respects. All right, you mentioned the products a little bit. Let's go back over them in a little more detail. What is the product portfolio for Splunk, the entire product portfolio?



MJ: Splunk's major product is the Splunk Enterprise server. That's our marquee product. That's the product that customers use to search, index, analyze their machine data. We have built a lot of applications on top of Splunk on Splunk Enterprise.

For example, we have a security app that's a premium app. We're doing really well in that segment. We recently announced a Splunk app for streams, and that enables companies to capture wire data as it goes across the networks.

Splunk Enterprise is our core product. We also announced Hunk. That is our product for enabling customers to do analytics on top of Hadoop. Hunk, essentially, is the same enterprise software stack that we took from Enterprise Splunk, and we deployed that on top of Hadoop.

CI: It's not just Hadoop, it's any NoSQL database. Right?

MJ: Yes. Hadoop and NoSQL databases.

CI: Yeah, interesting. Let's get into, "What does Splunk do?", because I think it's quite interesting. It's very different from traditional BI analytics. You do something, or you help companies with something, called operational intelligence.

First of all, why don't you define what you mean by operational intelligence? Then, secondly, if you don't mind, give me a quick example of something that would be operational intelligence.

MJ: Sure. In terms of the Splunk usage or what sets us different, is Splunk is actually a platform for analyzing your machine data. When we say analyzing, it's more real time. Splunk is not a platform that you'd use for batch analytics, it's for getting real time insights across your business.

What Splunk does really well, and why Splunk does really well, in most of customer environments is the fact that Splunk does not have any schema. Splunk's secret sauce is the ability for us to define schema on the fly, which means that as an organization or as a platform, we don't need to understand the underlying idea.



We can apply schema as companies search across this machine data, which means that, if your underlying data source changes, we don't have to change anything at all.

That's really the secret sauce, or the crux, of Splunk, and why we've been so successful. Secondly, we also have a universal indexer, which means that one indexer can be used for all different kind of data sources, which implies... Traditionally what used to happen was customers used to have different kind of indexes for different kinds of data. You don't have to do that anymore. You just use Splunk's universal indexer to forward data from all these different kind of data formats.

Totally, we don't have any RBMS in the back end, which means that we don't have to depend up on batch ETL to load data. The benefit of that is that there's no latency whatsoever, and, also, the fact that we are not defined by a rigid, back-end data base schema.

Lastly, we also enable companies. Because of our scale and performance, you don't need to filter in data. You don't need to sample data. You can get entire data, and you don't have extrapolations of the data, because you're using a sample data set.

In terms of "operational intelligence", what that really means is it's a journey that we take our customers on or our customers have been on.

I think the most common use case for our customers when they get started with Splunk is just the ability to search and investigate issues. My machine went down. Why did it go down? Let me investigate and try to understand what the issue was. That's the basic, or the first, use case that company gets started on.

But as they start using Splunk, and as they get more comfortable with Splunk, they realize that there is more value to be had from the platform. The next logical step for customer is to go into monitoring, so they can automatically monitor events rather than fitting for events to happen.

They can set up alters and thresholds and understand what's going on, what's going wrong in their organization so they can be more proactive to incidents, and they can respond to them in near real time.



The third use case, that we see a lot of customers using us for, is gaining end-to-end visibility across of entire IT infrastructure, essentially. That helps them with many cases. They want to make sure that from a KPI perspective, from a SLA perspective, their systems are up and running.

It also helps them, for example, make sure that they have the right capacity to meet the business demands. Lastly, I think where customers take the journey on is becoming the proactive organization by using Splunk for more real time insights.

Things that come to mind are examples in which companies that are using Splunk to understand how new device activations [are] happening, or how our customers [are] redeeming coupons on the websites. That's the journey that companies take from being reactive to proactive.

For us, the value of what we define “operational intelligence” to be is to gain insight, value, and intelligence from operational data that is coming from your machines.

CI: What's interesting to me is, as you go on this journey from reactive to more proactive, immediate, “let me act on something”, at least, in my work with clients and listening to others, it's far more valuable. The least valuable is the first one you mentioned. As you go up that kind of chain, it gets into you could save a customer from leaving. You could save a campaign, whatever it is. You get really good, valuable information the more it gets to be real time and the more they really push hard for that proactive analytic. Is that right?

MJ: Exactly. I think that's what customers have realized in using Splunk. And there are numerous use cases.

One example I can give is to talk about Domino's. The fact that they are the fourth or the fifth largest e-commerce website in the US is phenomenal, right? What they are trying to do is not only understand whether their systems are going to be up and running when customers are ordering pizzas with different channels, they also want to make sure, for example, how are coupons getting redeemed on the website. They want to understand what is the demographic of people redeeming these coupons, and they're



gaining real time insights from these coupons, so that they can be more effective in targeting their marketing campaigns in the future.

CI: Brilliant, it really is. Something else that we talked about a little bit was the business analytics, not operational analytics, but business analytics. Of course, we'll get into that I suppose when we talk about Hunk a little bit more, as well. What I liked was the position that Splunk has taken that business analytics is a compliment. It's complimentary, if you will, to the existing BI technologies like a data warehouse. Tell me a little more about that stance, if you will, because some companies are not saying that.

MJ: What we have realized in working with a lot of customers, customers have invested a lot of effort and time in building out data warehouse architectures, establishing EDWs, using BI tools.

You have to realize that 8, 10 years ago when I got involved in BI, I think there was a statistic done that said, on an average, a company had like six or seven BI tools. I think that era has past and you will see that more and more of these companies have standardized on their BI tool of choice.

That's because it's harder to introduce new concepts, new technologies, because it's a learning curve associated with that. I think keeping that in mind, I think what Splunk has done really well is understanding that it's beneficial for Splunk to have a complementary stance against all these different data sources.

What I mean with that is, a lot of customers have the machine data. They're using Splunk to analyze machine data. But they are using structured data to enrich that data.

An example being MetroPCS. They have partnerships for different wireless carriers across the US. They want to understand how cost effective their calls are. What they are doing is they are taking CDR records into Splunk and they are enriching that data with the cost structure.

The structured data contains information about cost, and they're trying to understand the cost of these phone calls that are being made across the network, and how they can optimize that better to reduce that cost, as well.



That's one example in which customers are taking advantage of the integration that we have with structure data sources. We have also have customers who are using Splunk to generate aggregates, and they're storing this aggregate in data warehouses they have built out.

Also, one more use case over here, there's our ODBC driver. We've had this ODBC driver around for almost a year. We announced a partnership with Tableau, around a year ago. What that allows customers to do is enable the non-technical users who are comfortable using Tableau to gain insights from machine data.

Earlier, it used to be very hard for Tableau users to get access to machine data, but with Splunk integration, they can easily get access to machine data. Not only that, they can mash the machine data that they get from Splunk with structure data sources, or an Excel file or something like that, so they can derive additional insights from the data sets that they have.

From a business analytics perspective, I think what we see, we see a lot of acceptance of us in four key use cases. It's digital marketing, product analytics, business process analytics, and customer experience.

CI: Let me ask you for one more case study, because I thought it was, these are marvelous, all of them are really good. But you also had one on the Internet of things that I thought was quite interesting. If you don't mind, just briefly talk about that one.

MJ: Sure. I think from an IoT perspective, one of the use cases that is pretty popular with our customers is a case study around New York Air Brakes.

It's also available on splunk.com, I'd encourage your listeners to go and check that case study out, as well.

But essentially, what New York Air Brakes does is develop brakes and these brakes are typically deployed on locomotives. So, any railway company in the US would deploy their brakes. First of all, they have sensors on these brakes and they're gathering data from these sensors and trying to analyze driving patterns and driving behaviors.



By capturing this information across different sets of railway companies, they're able to understand what leads to better driving, what braking mechanism leads to better driving and that eventually leads to more savings in terms of fuel costs.

Now, they are able to take this information to their customers and say, "Look, these are the driving patterns that are recommended for you folks, which will lead to better cost savings in terms of fuel costs."

CI: Yeah, that's just such a cool one. All right, I said we would talk about Hunk a little bit. We've got very little time left, unfortunately. But why don't you just briefly tell me a little bit about Hunk, a little more about Hunk.

MJ: Yeah, sure. Hunk is an exciting product that we announced, last year.

Unlike Splunk Enterprise, that is more real time, Hunk is a batch oriented analytic solution that is meant for Hadoop and NoSQL. Once customers have data at rest in Hadoop, they can leverage the capabilities from Hunk to visualize, analyze data that is at rest in Hadoop.

One thing to realize is that what we have done is essentially, we took the Enterprise product and we made some tweaks and changes to make it available on Hadoop itself. Just like the Enterprise product, we have a rich visual interface where customers can go and analyze data at rest in Hadoop. They can search query the data as well.

The unique things about our approach to analyzing data in Hadoop is that we don't move the data, the data stays within Hadoop, we just help customers search on the data set itself. The second thing, again, being a unique no schema architecture. We don't apply new schemas on the data in Hadoop. That's the initial benefit that customers gain from using Hunk.

We have customers like Vantix who are using Hunk to understand product analytics. They are provider of network gear that helps companies optimize network traffic. What they realized was in the usage of Hunk is trying to understand what kind of video streams customers are viewing more and what they are viewing less of.



The interesting insight they came out with was, if the bit rate of the video was much higher, then, customers would not be able to watch it all the way through. If the bit rate was lower, then the customers would watch those videos all the way through.

They can now take this information back to carriers like AT&T or any other wireless, any other communication providers and give this insight. Then, the communications providers can then take this information to do better capacity planning. To make sure that the customers have a better experience, and eventually, make sure that when it's time for customers to change, they do not do that because they've had a good experience with the provider.

CI: Yeah, boy, that's an interesting one. I love it. All right, unfortunately, we're out of time. I have to say that 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 Manish Jiandani of Splunk today. Thanks so much for speaking with me.

MJ: Thank you for having us.

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!