



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.bbbt.us.

Vendor:	Rocket Software
Date recorded:	April 17, 2015
Host:	Claudia Imhoff , President, BBT
Guest(s):	Bob Potter , Senior Vice President and General Manager Greg Willhoit , Managing Director & General Manager for Research
Run time:	00:15:54
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 my friend Bob Potter and Greg Wilhoit. Bob is the senior vice president and general manager, and Greg is the managing director and general manager for research and development for Rocket Software.

Welcome to you both.

Greg Wilhoit: Thanks, Claudia.

Bob Potter: Thank you, Claudia.

CI: Bob, let's start out with you. You said something that I thought was pretty cute, pretty funny. You said that Rocket Software is the biggest company that no one has heard of. Why don't you give us a bit of an overview of the company, and a little bit about its history, and some of the major milestones in the company.

BP: Sure, Claudia. In fact, yesterday, April 16th, was the 25th birthday of Rocket Software. Twenty-five years ago, our CEO and founder, Andy Eunice, started developing mainframe database tools out of his kitchen with his business partner, and offered those tools to IBM mainframe customers.

Eventually, IBM said, "Hey, we'd like to sell those under IBM brands," and Andy agreed to that. We expanded from one product to dozens of DB2 tools, and then IMF tools for the mainframe.

For many years, Rocket didn't have to have a market presence, because all the sales and marketing was done by IBM. Approximately 10 years ago, Rocket started to acquire companies, and we had brands that we continue to sell under those acquired names. Rocket started to rebrand



those products under Rocket, and we started building innovative products that were sold as Rocket branded products.

In the past two years, we've decided to really up the brand identity of Rocket, particularly over the last couple of years as we've innovated in the areas of data virtualization and business intelligence. I think now, with more of a market presence, attending very large trade shows and events, Rocket is beginning to become well known, but certainly \$335 million of revenues warrants that we have major awareness in the market.

CI: I agree. It's interesting. It is, in today's world, a relatively old company. Twenty-five years, you've had a good run rate there. I was quite interested in the fact that the company is keeping up with the times, if you will. It is involved in the newer innovations and initiatives like big data and no-SQL databases, even mobile technologies and cloud.

Why don't you talk a little bit about that?

BP: Let me start with database technology. As I said, we developed a lot of tooling for DB2, and DB2 on z/OS has some of the largest transactional applications in the world. We're talking about thousands of database tables with millions and millions of daily transactions, so we had experience there.

Then we made an acquisition about nine years ago, from IBM, in the MultiValue Database arena.

We bought U2 from IBM. MultiValue Database technology, even though it's been around for quite a few years, it's pretty innovative. U2 had thousands of ISPs and tens of thousands of customers, and we got a really good understanding of how other than SQL databases worked, application oriented databases.

A couple of years after that, we bought a mainframe database system called M204, which is a high-performance transactional database. It runs on System z. It's very innovative in its core technology, but we've taken those engineers and the U2 engineers, and we've decided to build additional innovative technology that we will be bringing market over the next couple of years.



You're going to hear a little bit more from my colleague, Greg Willhoit, about data virtualization, but we do have core database management skills here at Rocket.

CI: All right, let's bring Greg into the conversation, if you don't mind. Greg, let's do turn our attention now to the data virtualization product. Why don't you describe its architecture and features for me.

GW: Sure thing. The data virtualization product that Rocket had finished building and went GA last year is built primarily with the focus on data virtualization on z Systems from IBM, that is, on the IBM mainframe architecture.

However, it is not limited to mainframe data. It is basically built under the premise that the world needs two data virtualization architectures. To be effective, a data virtualization architecture needs to be able to be placed as close as possible to the data sources which it serves.

Therefore, with the plethora of disparate data types on the mainframe with the huge transaction volumes, the significant IoT-type data—Internet of Things-type data generated from the mainframe—we felt that because of the need for more and more operational analytics, the world needed that virtualization server with an architecture optimized for z close to the data.

We took advantage of many, many different hardware capabilities on System z as well as taking newer types of architectural paradigms like MapReduce and things like that, and bringing them to the mainframe, so bringing the new to the old, if you will.

In essence, we can support a plethora of data sources. We can support every single mainframe data source, IoT-type data such as Syslogs, and then SMF, or System Management Facility, which is kind of the original IoT. This is machine-type data generated by the mainframe and its devices, and connected devices and applications. There are literally thousands of virtual tables which we dynamically create, which allow for operational analytics against Internet of Things-type data.

Lastly, the architecture allows for a map once, that is, virtualized once, so no matter how many different types of application-consuming requests that you have, whether they're SQL, whether they're no-SQL for example...



You mentioned Mongo, Claudia... You only have to go through that mapping or identification of the data once, and then we simply expose it to all consumers.

What's even more interesting from a discovery point of view is every single Information of Things-type data that's generated on the mainframe, we ship out-of-the-box with the data virtualization product as fully mapped and fully virtualized, so there's no discovery needed. All the data is available. It can be then consumed by various products like Splunk, like Loggly, IBM Tools, you name it.

CI: Very interesting. There are a number of data virtualization tools on the market, and this was a question that we asked during the BBBT session. What is the biggest differentiator for your product, do you think?

GW: I think, Claudia, the major differentiation between the other data virtualization products and ours is that ours is built, is engineered, to run on z Systems. None of the other data virtualization products are. They do not run on z Systems. That's a significant differentiation right there.

Again, we simply realize, or we have this strong core belief the world needs two data virtualization architectures. The non-mainframe market is well served. I think there's some very good products out there. But, the major differentiation is we allow customers with mainframes to bring the analytics to the data, which allows for more real-time analytics as opposed to the alternative, ETL, taking the data to the analytics.

CI: Let's move then to where and why would an enterprise want to use data virtualization. Can you talk to me about some of the use cases and perhaps benefits of data virtualization versus the other types of data integration techniques?

GW: Certainly. We've seen many use cases where everybody's familiar with the nightly batch window, right? The nightly batch window has consistently been compressed. Year after year, customers have fewer and fewer hours and minutes to get their nightly workload done.

What that means is that they have fewer and fewer hours with which to move the data for analytics to be done off the mainframe. A major use



case for us, or for our customers of data virtualization is to eliminate that strain of the ETL load of pushing all that data, which is becoming more and more ballooned.

I mean, here's the problem: we're trying to violate the laws of physics here. We're getting to the point where the amount of data that we're moving is unwieldy, and so a major use case for us is to give customers a choice. We're not saying, "Don't move the data ever." That's not our message.

But a major use case for us is when customers are having trouble bringing their data into their analytics world, for whatever reason, because of timeliness, because of cost. A very nice use case for virtualization is to leave that data in place, pushing analytics towards the data, reduce the strain on your nightly window, and increase the responsiveness of your analytics.

CI: All right, Bob, let's turn to you again. There was a second product that you showed us. That was Rocket Discover, and I guess the same questions to you. What is its purpose, and what are some of its features?

BP: Thanks, Claudia.

Rocket Discover is a data discovery product that also has integrated self-service data preparation inside of it, or with it.

You know my background. You and I have known each other for a number of years. You know I worked at a data-integration software company, Expressor, which was acquired by QlikTech. I happen to think QlikView and QlikSense are very good products, but when I saw the data preparation vendors trying to provide add-on value to QlikView, I thought, "There's got to be a better way of having an organic, well integrated product that does data preparation and data discovery/data visualization on the same platform."

That was my primary idea.

I partnered with an 18-year Cognos veteran who runs the Discover lab, who also had two years of experience as a premier Tableau business partner. He also worked 18 months at SAS. His view was that the data discovery tools were not easy enough.



QlikView and Tableau, in particular, were still primarily geared toward BI developers who would build dashboard discovery applications and they were consumed by business end-users.

His view was that it still needed to be easier. We still wanted to reach a business end-user audience, and the combination of my data preparation thinking and his ease of use thinking brought us together to work on the Rocket Discover product, and that's what it is.

Those are the primary features of the product. It is architecture for the cloud. It's multi-tenant, browser-based architecture. You can run it in a private cloud, a public cloud, or on premise.

CI: And what I found interesting about the product was it has many different capabilities that I thought were quite good. It's got a little bit of collaborative capability in it. It certainly has the self-service look and feel. It just seemed like it had a number of things that were very useful and very modern in terms of its purpose. Would you agree?

BP: I would, yes. Its collaboration capability, I think, is the best in the industry. There's one or two other products that do it well, so I can't say it's thoroughly unique, but it's good.

We're unique in the fact that we can source mainframe data. Other products can't. You usually want to squeeze it through some other kind of translation technology.

But the real uniqueness of the product is in the self-service data preparation and the extraordinary ease of use that's in the product.

CI: Real quick, I understand it's in beta today. When will it be available?

BP: We have a dozen to a dozen and a half customers that are still running on the latest beta. We just did code freeze with Target Release 1 last week. The product will be launched and commercially available on April 30th, this month.

CI: Congratulations. In the remaining 30 seconds or so, who do you think the audience is going to be for this technology?



BP: The user type is a business user, a finance person, a marketing person, even a CEO or a CFO. But I want to say something about our go-to-market strategy real quick, Claudia. We don't have to sell to one single new customer. We have 15,000 customers already at Rocket. I have sister business areas and network management, backup and recovery. I've already mentioned database tools for ISD's.

We can completely focus on our existing customer base, and do very, very well with this product.

CI: That's an enviable position to be in, no doubt about that.

Unfortunately that's it for this edition of the BBBT podcast. Again, I'm Claudia Imhoff, and it's been such a pleasure to speak with Bob Potter and Greg Willhoit of Rocket Software today. Thanks so much to both of you for speaking with me.

GW: Thank you, Claudia.

BP: Thanks, Claudia.

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!