



[SAP HANA EFFECT](#)

*Title: Episode 9 - Bloomin' Brands Cookin' Up Innovation- (Duration: 25:41)*

*Publish Date: March 16, 2015*

*Description: James Williams explains how real-time planning and consolidation has [Bloomin' Brands](#) cookin' up innovation.*

THIS IS A RUSH TRANSCRIPT. THIS COPY MAY NOT BE IN ITS FINAL FORM AND MAY BE UPDATED.

**JEFF:**        **Welcome to The HANA Effect. I'm your host, Jeff Word, from SAP. Each week, we bring listeners the real stories of how companies are taking advantage of real-time computing to transform their organizations and let them share the lessons they've learned along their journey. Welcome to another episode of The HANA Effect. I'm Jeff Word, from SAP, and I'm here today with James Williams from Bloomin' Brands. Welcome, James.**

**JAMES:**    Oh, thank you.

**JEFF:**        **Let's for a second talk about what Bloomin' Brands is.**

**JAMES:**    Bloomin' Brands is not recognized as widely as what we do. Bloomin' Brands is the parent company for restaurants: Outback, Carraba's, Bonefish, and Fleming's

Steakhouse. We have an international presence and our home office and corporate headquarters is Tampa, Florida.

**JEFF:** **So, I'm guessing the Bloomin' is from Bloomin' Onions in Outback, right?**

**JAMES:** Our most world-renowned product.

**JEFF:** **One of my favorites, by the way. So, you guys are a phenomenal SAP customer. We're thrilled to have you guys here talking about your HANA experiences. So, what's your role at the company, James?**

**JAMES:** I lead the SAP Basis team and the database administration group, which includes HANA and Oracle and SQL Server and Hadoop.

**JEFF:** **So, you are actually responsible for making sure all of this HANA stuff actually works so that the business can do the jobs that they need to do?**

**JAMES:** I'm the person who translates what SAP does. It's actually been a really, really good product for us. It's functioned well, and to be honest, it just makes me look good all the time. We have a team of users who are basically subject area experts in their domains and they are power users with respect to querying data from our data warehouse and from our financial warehouse, which is now HANA. And I just keep the system up, and the users are doing a

fantastic job of answering the businesses questions in a great deal more haste than ever before.

**JEFF:** **Well, that's phenomenal. I love hearing stories like that. Let's take for a second and explain to people kind of why you guys went down the HANA road. What really was the beginning of this discussion and the project for HANA? What really drove you guys to consider that and start down that path?**

**JAMES:** We started down the path because we had made a decision in our adoption of SAP to have business warehouse on HANA. We looked around and we are an Oracle shop and we have Oracle and Exadata. But we wanted to have some financial close goals. We wanted to be able to do business planning and consolidation with a great deal more rapidity than we could do and then we did before go live in HANA looked like the right thing. We knew other customers were using the business accelerator, which is the precursor or HANA type product and, thereby, HANA made sense. While HANA was very new and we were very new to SAP, the expected ROI was exciting. So, we went down that road.

**JEFF:** **It's fair to say that performance was really one of the key things. You needed to do your financial closing much faster than it was being done on your existing database. Is that fair to say?**

**JAMES:** That's exactly right.

**JEFF:** **So, that's typical, actually we hear that from a lot of customers. The speed is the sexy part that gets everybody kind of interested. What would you say the kind of reaction from the team was once you started to really get your hands on HANA beyond the speed?**

**JAMES:** It's actually a new world. So, the good thing is that the team was ready to embrace new tools and to begin to look at data more smartly. I mean the whole columnar approach was something new, especially in business warehouse where the folks had become used to looking at some of the data cubes in a certain way. And now, HANA was making that process so much faster. But, again, it was our standard processes. So, while I want to get away from talking about the ultimate performance, but even from a business perspective when I have the financial planners who normally would load data and run a planning model in 4-6 hours and now they're doing it in 10-20 minutes. Their scenarios changed. There are a number of planning cycles they can run through within a single close is just remarkable given what they had. So, the quality of the information that they're passing to our executives is so much higher due to the ability to run consolidation models and planning models so much faster in HANA to have that turnaround. I keep coming back to speed because speed is the sexy part, as you said, but all of the business functions

that we're accomplishing the goals are the key. But they're happening now in minutes and seconds instead of hours.

**JEFF:** **That's a phenomenal point and something I think might be lost on some companies that you guys have found in spades is that the speed is an enabler to do more important and more expanded and more smart types of things with the data. Without the speed, you couldn't do those things. And it's really great to hear that you guys are getting that value. Let's discuss for a second what really happened when you guys started the project? Can you tell us a little bit about the structure of the team and what you guys did and some of the big steps that you went through to get up and running on HANA?**

**JAMES:** Well, to be honest with you, we had a very small team. And I think that helped us I that we brought in a bunch of our people who knew database very well. And, initially, a part of the HANA adoption was to say, ah, it's nothing new. It's just database. We know database and we're going to get this done. And that actually carried us a very long way. So, we were able to get data cut over. We were able to build our connections through SLT for our ECC environment and we were able to build a very robust HANA environment very quickly because we didn't acknowledge HANA and HANA let us do that. We used SQL Server experience and Oracle experience to deliver HANA results. It was only after we built that that we wanted to refine and

wanted to get greater performance or we were just in disbelief at the amount of compression we were actually getting, for example, in HANA. We were saying, wait a minute; this data couldn't have loaded correctly because if we take a terabyte and we wind up with 300 gigabytes, well, something's wrong. So, going through those things, HANA was performing well and it actually did a bigger job for us. It allowed us to develop with it. You know, when you look at a team, and when our project started, current HANA release was release 53. And when I went off to HANA class to get my certification or get my training prior to certification was release 58. So, I was an Oracle guy. And I translated everything that I learned about HANA into how would Oracle do this? HANA let me do that until I got proficient enough that I could actually discern the differences that existed and the decision that needed to be made that were truly HANA centric. From our project size, if we had a large project, I think it would've been harder to divide up the work and the learning across a larger team. Just the three of us actually who worked on this thing, we were able to dissect it and do just what we needed to, one, deliver to our customers and, two, to learn the emerging HANA environment. From the start of our project at HANA 53 until the time to go live where we were 59.9, there was a lot of learning they had to do and a lot of change in the product as well. The product matured while we were implementing. It's kind of like implementing on a treadmill, right. But I appreciate the fact that we were able to always have documentation prior to needing it. We

had issues and we called up and we could get knowledgeable support from people wherever they were. But there was always good, solid documentation to cover whatever feature or release we were planning to go to or looking at going to or on. Again, those things were a big part of us being able to have a small team during this implementation.

**JEFF:** **That’s phenomenal and we see that again across the board. But you brought up something very, very important that I want to take a little deeper discussion into. You talked about that journey as a DBA or a database guy going from the Oracle or the Microsoft world to the HANA world. You mentioned something very, very cool because the first kind of portion of that process, of that evolution that you went through, you would take HANA stuff and concepts in HANA and then kind of reconcile that with the way things were done in the world. But at a certain point, it sounds like you flipped a switch and started thinking the HANA way. Can you talk a little bit about some of the things that were going through while you guys were doing this that really triggered that shift from the old world thinking to the new world thinking?**

**JAMES:** The first bit was truly embracing what it means to be column store vs. rule store. That’s where we made some mistakes early on. But some of our philosophies changed

as we matured in to understanding HANA and understanding what column store is. Just looking at the table, we have some key tables in HANA ECC that are very wide tables, 300 columns. And in the Oracle way, we would look at this and sometimes put views or otherwise we'd just have to endure the performance. Well, now, when we start to understand the power of columnar, well, we just list out those columns that we need as you might expect in a normal SQL statement. But the HANA's engine is just going to deliver those columns, and I mean what you're moving through your in memory database and what you're traversing is so much smaller. And you learn things about compression and you come to this with a sense of database. And you say, well, I know what a block is and I know how much data is in a block. Well, now, when you look at this compressed HANA page and the amount of data on your page that's being brought back, I mean it only accentuates how much speed you're going to get from an in memory database at all. I'm perhaps being a little bit disorganized in my thoughts, but embracing all those things because they're all out there and they're all being really talked about right now. But with HANA we have a representation, something we can really work with. And as we start looking at the importance of in our SQL statements requesting the columns that we want. I mean one of the rules of thumb in the HANA database, you never say "select \*." We have developers who tried that. Who said, "Wow. It's much faster than naming the columns we want and we could demonstrate how much additional labor

we were putting on our in memory database when they did that batching and how robust the database when we only requested in a query the data we needed.” It changes your thinking. It changes the way your development standards are made in your application organization. It changes everything and it comes over time. You learn what’s poor practice, and you move into best practice really, really quickly because poor results and good results are seconds away at every turn.

**JEFF:** **Putting yourself in the position of a lot of the other DBAs out there in the SAP world, what kind of advice would you have for them when they’re starting to think about their SAP project and the philosophy that they need to approach those new projects for HANA with?**

**JAMES:** As we bring on new DBAs and we take people who hadn’t been involved in the implementation, the first thing I do is make sure they have a good, robust understanding of what we really mean when we talk about column store. And what that really looks like in the database and what the database retrieves when we’re talking about column store. Again, right now, there’s so much information out there. But when you understand that, you really begin to get the foundation of HANA understanding. If you understand the differences between column store and row store and when they should each be used, it’s going to be probably our biggest design tool. There will probably initially be a

reason to think and to design in a row store centric manner. Don't do that. Figure out what the penalties are. Figure out what you really need to accomplish and learn how to use that. When you look at your views, coming from an Oracle perspective, a view is a view, and sometimes we have materialized views. But, of course, in HANA we now have aggregate views and calculation views and so forth. So, when we look at some of our financial data and just some of our analytical metrics stuff, the kind of marketing kind of information, we don't have to do roll ups and summaries and aggregates anymore. This is accomplished at the database engine in HANA. That alone is a pure development philosophy change, but when it's embraced, everything changes. You have a really hugely powerful information platform. In our other database environments, we often spend hours at night building aggregates and rolling up data so that it could be simply reported on by our tools, such as business objects or OBIEE. HANA accomplishes that natively with its views, with its special views, the calculation views and analytical views that I mentioned earlier. With this, there's none of that and we are right about producing informational results without the hours of data gathering and consolidation that formerly took place. The speed of the HANA engine is one thing, but the hours that are given back if you look at our business where we have results that come when our restaurants do business and when they close and we want to know what our results might look like at the time of a close. We don't have to take that data, capture it from all

of our stores, process it overnight and have our business model look at yesterday's data today every day because it takes so long to process it. We can see a night's results right away.

**JEFF:** **That's pretty amazing. You can actually watch people paying their bills at all of your stores in real-time, huh?**

**JAMES:** Almost.

**JEFF:** **Let's talk for a second about some of those outcomes. You talked about changing this overnight model, changing these 4-6 hour runs down into a few minutes. What have been the business impacts from that? What's been the really big results that the business users have seen from having his new, powerful engine from HANA?**

**JAMES:** It's interesting because they get to use the results better and they get to do what they're really here to do. They get to actually use the information and do what-if scenarios. So, not giving away any business secrets but if a person says, you know, hey we sold a million onions last night. What if we discounted those onions by 39 cents? What would the ultimate impact to cost of goods and so forth be? That kind of thing really would take days to, one, collect enough information from the right places and then build all the aggregates that I talked about earlier. But it's

done now just in the course of a session where a user sits down and just plugs in what they want to see and can introduce variance of any kind and get new results, so that they get to think about how to make our products more cost effective to the market or how to determine if a product just doesn't make sense anymore, just instant gets results. And that changes things. It's kind of almost passé talk about an in memory database being fast. But that's what it does; it allows people to look at information and to process it at human thought speed. It's far more productive and it's resulting in innovation in our business.

**JEFF:** **Those sound like some absolutely amazing results and huge business impact. What were some of the surprising things you guys discovered during the project that you weren't expecting?**

**JAMES:** So, as I approached this, we were really open and were really positive about HANA in the beginning. So, we didn't have a great deal of expectation outside of, wow, this ought to be fast. But we were really surprised to learn how fast things were. One of the tables that are really heavily used has 295 million rows in it as of a couple of days ago when I last looked at it. When we ran certain queries against that on our older high powered platforms, we would have queries that would run 8-9 minutes. And to get that query to run 8 or 9 minutes, we actually have DBAs spending a lot of time doing maintenance against those objects to get them to work well. You have that table in

HANA and we actually without exaggeration process the same queries and they come back in something like 29 milliseconds. That's surprising. We expect it fast; we didn't expect how fast. The other surprise that then came is when we adopted HANA, all of the information be it marketing information or technical information led us to believe that we should expect a compression rate of four to five times. What we're seeing is five to 11 times and sometimes even greater than that. The reduction in the physical size of the data decompression it just means that we're moving more data with every iOp that we execute. So, that was pretty impressive and a great surprise to us. Because of that, we've adopted more use of HANA because we wanted to look at more things; we wanted to do more processing in HANA. We had this high compression rate where we wanted to pass data out of HANA into legacy databases and legacy systems. Those systems began to grow in at an extraordinary rate because of the large volume of data that existed in HANA and decompressed, if you will, as we extracted it using tools such as data services to import them into Oracle databases or SQL Server databases. It became practical to actually leave the data in HANA. So, our HANA growth in the adoption of systems that we had no intent of bringing to HANA was surprising. And it came because it was faster and, believe it or not, cheaper to keep the data in HANA.

**JEFF: Well, that's a very important point. The compression ratios people tend to kind of underestimate them,**

**like you guys did. And they're usually surprised with that. But it's the ancillary effects of the smart way that we store the data inside the columnar databases that translates into a much lower TCO across the landscape rather than just looking at things in just the HANA box. So, you guys have had phenomenal success with HANA in your data warehousing strategy. Can you give us a little bit more of an idea of where things are going next? Based on this success, where do you guys think you'll be headed for your next few steps with HANA?**

**JAMES:** One of the things that we are really excited about is how HANA can help us build a 21<sup>st</sup> Century loyalty program across our brands. We want to know, as immediately as possible, when a person has been at our stores and what products they order and maybe we can offer them some creative discounts or greater hospitality by knowing when they were at our last store and what their orders are like, and so HANA is the only tool that we have that could give us the kind of rapidity to process the data to tell you, "Thanks for coming in, Mr. Word. Last time you were here, you had a Bloomin' Onion and a Coke. Can we start you with that? I mean that kind of speed of information and to use that to offer greater customer service and to build brand preference is where we're looking at. And that's what we like about what HANA offers today.

**JEFF:** So, one of the things we're getting lots of questions from people is, could you tell us a little bit about the hardware platform and the landscape that you guys decided to deploy HANA on? Just so that people get an idea of what's the hardware and infrastructure you use to make this HANA system work.

**JAMES:** With respect to a HANA platform, we started out with a very robust platform. We started with a 2+1 scale out environment. Scale out is a term used to talk when we are actually merging the processing power of multiple physical HANA servers into one single logical environment. And I would say that the biggest part of being successful with HANA from a hardware platform perspective, it decides what you need. If you have a database that's 100GB in size in your legacy platform, don't buy a 2 terabyte HANA machine. You're not going to grow into it. With compression and with actual data speeds that you're going to achieve, you're never going to grow fast enough to make use of 2 terabytes of machine space. The machine will antique before you ever get to use it. So, pay attention to the HANA sizing models. There's a quick sizer tool that you can use. Those hours that you spend agonizing over the size of the platform will be dollars returned to you in ROI as quickly as you get a productive system in HANA. Our company started with a multi-node machine. And we determined that having multiple single node machines is cheaper and more efficient than we'd ever hoped. A scale out environment that had 2+1 means

that there were two operational nodes and one that set in standby awaiting the failure of one of the two operational nodes. Because they were 512 gigabyte machines, we effectively had a logical 1 terabyte machine with 512 sitting, waiting for a failure. Unless you have a database that is larger than the physical size of the machine you're buying, thereby, if you have a 512 gigabyte machine, that database should be bigger than 512 gigabytes or else a scale out machine doesn't make a lot of sense. And it'll be hard to really get the financial returns that you want. Again, we moved away from our scale out machine and implemented multiple single node machines. We've had a great success with that and we want to continue down that road.

**JEFF: So, that's great advice and everyone would be very well served to spend some time looking at the SAP notes on the sizing and going through the quick sizer like you suggested. We recommend to everybody, take a look at that before they embark on the procurement process for their hardware. And it's also very good advice to really stop and think about what your needs are rather than just using the traditional methods and metrics for sizing hardware. HANA is a different beast and you have to treat it differently like that. And so, I'm glad to see you guys are on the right track and happy with your system now, especially since it's going to become more and more important to the success of Bloomin' Brands**

**over the next several years. With that, I want to say thank you so much for taking the time to go through this with us.**

**JAMES:** Well, I'm pleased to have helped and I'd like to be a part of anything going forward. We are happy HANA users and it's in our future. And the next time you go to a Bloomin' Brands restaurant, think HANA.

**JEFF:** **I will absolutely do that when I'm ordering my next Bloomin' Onion. So, with that, I want to say thanks again, James, and to the entire team at Bloomin' Brands for being on the podcast today and for being such a wonderful SAP customer. If you guys would like more information about the Bloomin' Brands use case or the different topics that were discussed today, please visit us at our blog at [hana.sap.com](http://hana.sap.com). As always, you can follow us on Twitter @hanaeffect, and please subscribe to us on iTunes so that you can get up to date every week, a new episode of The HANA Effect. And with that, I want to say thank you and "Tschus."**