



*SAP HANA EFFECT*

Title: Episode 1 - [Jabil Circuits Rides the Sidecar](#) (Duration: 20:17)

Publish Date: January 20, 2015

Description: Paul Jackson, ABAP Manager for [Jabil Circuits](#) stops by to explain his company's first HANA project: ERP Sidecar to accelerate ABAP inventory reports. Listen how Paul's team was able to reduce reporting times from 8-10 hours down to 3 minutes using SAP HANA as a sidecar database— with a part-time team of 3 people and no budget!

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 their lessons they've learned along their journey... I'm here this week with Paul Jackson, the ABAP Manager from Jabil Circuit, a great SAP customer. Thanks for coming, Paul.

**PAUL:** Thank you.

**JEFF:** Before we get started, tell us a little bit about what Jabil is and what you guys do.

**PAUL:** Sure. Jabil is a global contract manufacturing company. We specialize in electronic goods operating pretty much globally 24/7 around the world.

**JEFF:** **Paul's going to be talking today a little bit about the use case that you guys have used HANA for and it's one of the more common ones. Kind of the early starting points for a lot of our customers, and that's HANA as a sidecar. Basically, for those of you who aren't familiar, you take a small HANA server; you stick it next to your production database underneath your regular old you know SAP ERP system and what you do is you set it side by side. You copy some of the tables over to the HANA one and offload some of the heavier duty queries and workload from the production system onto this new turbo-charged HANA system. So, Paul, first of all before we get into what you guys did, what were some of the burning business issues that your company had with the SAP system at that time? Why did you have to kinda start thinking about a different way to solve these problems?**

**PAUL:** Mainly volume of data together with being a global company that really operates 24 hours and seven days a week, pretty close to. So, you know, the big data, the huge MRP runs, we have processes that if that delayed an hour, production can go down. You know, so the slightest hiccup in production affects a five hour MRP run that

affects everything downstream from there. And then, you know, we get people yelling at us because the factory floor is down, essentially. Or, the suppliers can't pitch us or whatever. So, mainly it's -those two points—the fact we operate globally with no real downtime and the amount of data that we have in our system. We have I think it's somewhere around 30 terabytes in our database.

**JEFF:**        **Underneath your SAP ERP system?**

**PAUL:**        Right, yeah.

**JEFF:**        **So, that's a pretty important point that a small or even a large amount of latency in the system can have some serious business impacts as it snowballs down towards the production line. So, performance to you guys is a necessity, right? So, what was the system kind of some of these queries, what was the situation before you put HANA in?**

**PAUL:**        One of our big pain points was the set queries that our user community was creating. - - you know massively inefficient joins particularly in warehouse management. - at times, they would bring our database to its knees almost. We'd have to get - - we would get into kill jobs and that sorta thing. So, you know, it was a real drain on our system resources. So, we really felt that that was a good place for us to start.

**JEFF:** That's great. So, what was the big impetus for HANA? What made you guys start going down that path of potentially using HANA to solve these big latency issues in the database underneath your ERP system?

**PAUL:** To be honest, I'm not, I really wasn't involved in those decisions. From my perspective, it kind of just appeared on my doorstep one day.

**JEFF:** So, somebody from up above said, "Let's give this HANA thing a try and see if it can fix this problem?"

**PAUL:** Yeah, pretty much. I mean, you know, we have such huge problems with the size of our data that anything that looks like it's gonna help us, we're gonna take a look at.

**JEFF:** So, walk us through kind of what happened once it landed on your desk and somebody said, you know, give this thing a run. What did you guys do? How did you kind of proceed from that point?

**PAUL:** First thing we did is played around in a sandbox system just to figure out exactly what we could do with sidecar. So, yeah, initial thoughts were that using traditional web app programs that we would just connect to the HANA sidecar as an external database, which really required very little code changes to existing reports. It's really the connection keyword and just select the database

connection that's been configured. So, from sandbox we played around. We, obviously, thought we saw some impressive results. But, you know, the data in our sandbox systems is very small. So, we really just in terms of syntax and what we could do, we just trialed it in our sandbox system. And then once we confirmed that this is something that - - could really help us, at that point, we got together with some of the ECMAs and the functional managers to determine where the best place to start would be if we were take this into a productive system. What would be the low hanging fruit, essentially?

**JEFF:**        **What were kind of some of the general reactions to all the newbies that were just coming into this HANA project? I mean it was your first time, right? So,**

**PAUL:**        Yeah.

**JEFF:**        **What did people start to think once they got to play with it?**

**PAUL:**        Well, we were excited cuz, obviously, we'd seen, you know, a number of us had gone to TechEd. I think some people to Sapphire. I don't get to go to Sapphire. So, I don't know what they talked about there.

**JEFF:**        **It's a lot of drinking and parties. That's about it.**

**PAUL:** That's what I hear. Yeah. And so there was a general excitement anyway. You know, we knew what was being talked about in terms of what HANA.

**JEFF:** **The potential yeah**

**PAUL:** The potential, right.

**JEFF:** **But what happened once you actually saw the realization of that in these tests in your pilot?**

**PAUL:** That's when we really sat down and realized what an impact this could make, initially, to some of the queries I was talking about. And we still haven't really gone beyond that point, although we plan to. But, obviously, you sit down and you start looking at other things that this could affect that could really make a difference to us as a business not just as an IT - - what are we?

**JEFF:** **Group.**

**PAUL:** Group, yeah. That's it.

**JEFF:** **So, that's a great point. You started off trying to solve one specific business problem that you're having with system performance and latency...**

**PAUL:** Mmm hmm.

**JEFF:** **But then once you started playing with it, sounds like things started to kind of snowball from there and people kinda jumped on board**

**PAUL:** Yeah. We actually had one group. So, the IMWM team really were the ones that fronted up. Partially because of like issues with how inefficient a lot of the queries were. So, they really rocked up and put their hands up and kinda said, this is something that we see huge improvements. And, of course, everybody was on board from Netweaver because these queries are such a headache, ABAP because you know it was gonna be a easy way to trial what we had in IMWM. Because they're gonna get the end result. And so, from there we worked with one of the IMWM functional specialists. And that's where we sat down and mapped out two main areas that we could improve upon. The first one was the inventory movement. And the second one was inventory aging.

**JEFF:** **While you guys were going through this, HANA's a little bit different environment than most DBAs are used to. It's exciting. It's cool. It's new. But it's not the same old same old. So, were there any type of interesting kind of realizations that people had during this testing process that they were surprised by something or something that they didn't expect happened?**

**PAUL:** Not really. I mean there was a little bit of trial and error. But, you know, using the sidecar approach how we started with the - - initial programs that we developed weren't very different from the ABAP programs we traditionally had for the same job. So, you know, we encountered some restrictions that we noted. You know, cluster tables, for instance, were a problem. You can't query cluster tables using the connection keyword. So, we had to build SQL strings, essentially, to query cluster tables. But, everything else was similar. The slightly different mindset in that we traditionally, we have limitations on how many tables we like to join and what type of tables they are. And we kinda threw that out the window and just tried to join as much as we could all in one query.

**JEFF:** **How'd that turn out for you?**

**PAUL:** Turned out pretty good. I think some of our initial code reviewers scratched their heads for a while, but once we got the message through that, you know, that's gonna be acceptable on a HANA database, then we didn't have any issues with that.

**JEFF:** **That's actually a great thing. The new rules that programmers need to kind of adopt and live by are very different than the rules they grew up with. So, you kinda have to flip that switch in your head when working with HANA to really take advantage of this new paradigm, right?**

**PAUL:** Yeah, it's true. You have to almost reeducate people to say what used to be frowned upon, now might actually be better, you know. There's other areas we struggled with too. There was, one of our first production versions of the inventory aging report I think it was. We tried to do an aggregate function inside a loop. We were interested in how that would world. And that didn't work out very well. So, I mean, it was still fast. But we looked at why it wasn't faster. It initially surprised us why it wasn't faster, and we determined that actually most of the run time was being spent in this aggregate query. So, we don't like to use aggregate queries in a loop.

**JEFF:** **So, slightly new rules of gravity to play by?**

**PAUL:** Yeah.

**JEFF:** **So, let's talk about the outcomes of this. So, what happened with you let the users know, okay, it's ready. Let's do it. What was kind of their reaction once this thing went live that they could use on a daily basis?**

**PAUL:** They were very impressed, obviously. Initially, they were resistant because we were taking away the set queries that they'd been using for years. But, you know, when you compare a 12-hour run time with a 3-minute run time, then they quickly figured about the queries they were using. And it took a little while. I think, you know, we

rolled it out and we made the user community aware that it was available. Obviously, those people who lost their set queries needed to know immediately. Others we just informed. Because we were also interested what else we would uncover if we made this report freely available to certain groups of users, who would come outta the woodwork basically and start realizing that either they can use the report as it is for something else they were doing, or they could or request a similar report with just a few changes. And that's actually what's happened is the longer it's used, the more request we're getting to add new data. Calculate things in a different way. So, we replaced 250 queries with one report in the case of the inventory movements. So, that's a big report, right? We ended up 120 something columns, which we manipulated using just local and global layouts for the users. So, the users don't really even know there's this huge thing behind the report they're looking at. To them it's the same 30 columns that they had with their query. And then, of course, as soon as they see it's a report that can be run in three minutes, they start thinking, oh, well what if I added this in? It would that would enable me to be able to do this other part of my job too. And I'd be able to get the results in three minutes for that as well. And so, we've had this endless flood of requests for mainly for changes to the existing reports but also a couple of new repots. We've branched out into some SD reports now. The main thing holding us back right now is funding to really ramp this whole thing up. But, the results can't be disputed. And I

believe we're close to getting that fund now so that we can really get some people onto it and really start making a big difference in other areas as well.

**JEFF:** So it sounds like your original problem was having these really slow reports that people couldn't use very often and the data was kind of locked up. Right? So, you used HANA to unlock the data, make it easier and more accessible. And, now, people are just feasting on this data all over the place and requesting more and more and more of it because it's so easy to get their hands on.

**PAUL:** Right. That's exactly what's happening.

**JEFF:** If that's kind of the case now and you're starting to get funding, what's next? Now that you've got this momentum, now that you've got these users that have seen the power and the value that HANA's gonna bring to their daily jobs, you're reducing all of this potential down time of the production lines and things like that. Where you guys gonna take it to the next level? What's the next big steps for this once you get the funding and resources?

**PAUL:** That's a very good question. I'd like to know the answer to that too.

I honestly don't know the answer to that. I mean we need to take it beyond what we've done. We've really just been playing around at the moment. You know, in terms of where we would like to go, the really big hitter for us would be MRP. If we got our MRP onto HANA, that would be huge. That would be a game changer for our company. You know, we can continue to look at the other queries and start replacing them. But, there's obviously value in that, but I don't think that's gonna be the really big payoff for this. That's kinda just where we started with our little playground to see what we could do. And there's a lot more we can do than what we have. That's for sure.

**JEFF:** **So, let me ask you some couple questions about the team that did this. So, you mentioned that you had lots of people from the business that were involved. Lots of the users -from inventory and warehouse management. What was the technical team like that you did this? How many people? What were their backgrounds?**

**PAUL:** Very small team. This was really done I wouldn't say in our spare time but, certainly, as a peripheral exercise to the other tasks that we had. So, from a Netweaver perspective we do have some very good people who, in terms of setting up the sidecar and the SOT replicator and all of that, they do spend a lot more time on it than development, the ABAP development team did. And they still do. But, from a ABAP perspective, we worked with just

a single functional analyst on the IMWM side. She communicated to the user community. I did some initial development myself in sandbox. And once we proved out what we could do, I had one of my developers take it from there, again, just in his spare time, I guess. So, we allocated like you know 10 maybe 20% of his time. And this just ticked took over for a while, which is part of my problem now. It's become so popular, I need to find someone to put onto it at 100%, which comes back to the funding. But it's kinda become something that's snowballing so much that I can't just assign the odd developer at 20% to it anymore

**JEFF:** **So, it's becoming mission critical now, right?**

**PAUL:** I wouldn't say mission critical yet because they're still just inventory reports. It's not anything critical to the business at this point. But, in terms of generating excitement from within the organization at every level really, then it's definitely doing that. You know, so we can now move into something that's really gonna make a difference to our business. And people are prepared to fund that now. Whereas before they weren't prepared to write the big checks, I think they're more prepared to do that now that they've seen what it can do and that it works.

**JEFF:** **Yeah, so you guys are a great example, and we see this across our customers around the world, trying that first step, that first experiment to see if it**

**works. You guys did it part time with you know, two or people, right? But once it worked and people saw it and saw the value of it and saw this, you know, massive improvement in the process, that starts to snowball. And that success kinda breeds more success and more engagement from the senior people with the big checkbooks, like you said. That seems to be the pattern that most of our companies are following right now. The great thing is, at some point, you guys are gonna have to think about this, right? So, tell me about that. What do you guys think about instead of going sidecar, which is very non-disruptive and very low risk, great way to start, what about the big enchilada down the road? Putting the suite on HANA; putting the ERP entirely on HANA. What have you guys thought about that? Is that part of your roadmap?**

**PAUL:** It's been discussed. I'm not sure where it is strategically. There's a lotta discussions that I'm not privy to. I know that it's on our roadmap, but I'm not sure that we have a definite timeframe around there. I think, at this point, as I see, we've got BW Enterprise HANA that seems to be going well. I think we've been on that for over 12 months. At this point, we're looking at ramping up to sidecar. As I see it, it's on our roadmap but I'm not sure that there's a timeframe around it.

**JEFF:** No problem. Couple of extra questions. So, you guys did this pretty much entirely in house, right? You didn't have any external consultants come in?

**PAUL:** Correct, yeah, in house.

**JEFF:** And you did it part time with just two or three people and made a huge, huge business impact. So, in general, -do you think that companies need to develop that type a skill set themselves and do these pilots, you know, on their own or is it something that you think that it might be different enough that they need to bring some outside experts in?

**PAUL:** No, I mean I think in terms of ABAP development for the sidecar using the approach we took, then you don't need any special skills. You know, what you already know is sufficient to start using the sidecar. And that's the beauty of it, really, is that we didn't need to go on any special training courses. We just used our existing ABAP knowledge to redirect queries onto a different database, you know. And then once we got there, we played around. We figured out how best to use it. I mean, one of the things we did was to create like a utility cluster connect to HANA. So, we put some standards around that to insure that everybody was using HANA connects to it using this utility. And, you know, that adds additional functionality like insuring the connection's available, and if not default them back to Oracle and communicating with the HANA

DBAs and by our email that the connection's down and that sort of thing. So that just puts safeguards around the fact that you know, we have a backup. For whatever reason, any communication issues with the HANA database, then we have the Oracle database to fall back on saying that the code we're running against the HANA database is not really optimized for the Oracle database, but it does at least handles errors in a nice, seamless way without any nasty, you know, terminations.

**JEFF:** **It's a nice fallback to have.**

**PAUL:** Exactly.

**JEFF:** **Status quo is a nice fallback to have, right?**

**PAUL:** Exactly. It makes it makes it a little like we know what we're doing when users don't get terminations and short dumps in their programs.

**JEFF:** **Well, Paul, thank you so much for sharing your story with us. Thanks for being here.**

**PAUL:** You're welcome. Thank you for having me.

**JEFF:** **Again, the goal of this whole exercise and podcast series is to share those real world stories from guys like you that, you know, have blood on their hands from the projects and really had to live through this.**

**So, for that we really thank you. And, on behalf of the listeners, thank you for sharing that with us. If you guys want to continue to hear more episodes like this and more stories from great SAP customers using SAP HANA, go ahead and subscribe to us on SoundCloud and iTunes. For more information about the use case that Paul just talked about, go to [saphana.com](http://saphana.com). You'll find lots of great information about the sidecar use cases. In addition, if you're an ABAP'er out there and you want to know more information about SAP HANA and ABAP, my great buddy, the world famous Thomas Young, is teaching an online, free MOOK, on ABAP on HANA. So, go to [saphana.com](http://saphana.com) and you can find out about that. And, as always, to keep up for updates, follow us on Twitter - - @hanaeffect. And with that, we'll say thanks and "Tschus". [END]**