

Blog > The “big ideas” of today really are mostly very late realisations of the ideas from half a century ago.

The “big ideas” of today really are mostly very late realisations of the ideas from half a century ago.



Journey Through SAP
with JZ

Lars Breddemann
Senior Database and Analytics Expert



The “big ideas” of today really are mostly very late realisations of the ideas from





Jarosław Ziółkowski
Sales Executive

 2020-10-09

[MACHINE LEARNING](#) [SAP HANA](#)

Let me introduce my guest

My today’s interlocutor transfers customers’ thoughts into the source code and has been involved in the SAP industry for over 15 years. He specialises in database-related matters and loves to solve unsolvable problems.

About Machine Learning in the SAP industry, the differences between On-premise & Cloud solutions, and his adventure with SAP PRESS Book – Lars Breddemann is interviewed by Jarosław JZ Ziółkowski.

Reading time: 9 minutes

1. You work as a senior data management and analytics expert. Thus, I need to ask: what exactly your job is?

In short, my job is to know how to work out what the client wants and needs and how to implement those requirements. Often, I get pulled into an implementation project where there already exist data and code, and the question is how to make the process faster or more scalable.

The thinking here is “oh, we just need a little HANA-magic-tuning-dust to be sprinkled over the solution”. That, sadly, is not how things work in real life. When I work on a solution, I find the bottleneck and the reason for why this bottleneck

Search:



Related articles

Smart Data Integration (SDI)
- transporting SAP HANA objects

Uladimir Sakalenka
SAP Technical Consultant at Int4





[int4](#)

I want it all and I want it now! Smart Data Integration (SDI) – Initial Load and Real-time Replication

Smart Data Integration (SDI)
- transporting SAP HANA objects

Uladimir Sakalenka
SAP Technical Consultant at Int4





[int4](#)

Smart Data Integration (SDI) – transporting SAP HANA objects



For example, I had to optimise a big financial analytics process and found that certain “master data” lookup values were re-computed every time the process was executed. By changing the solution to compute and store the values upon data loading time, I was able to remove this part of the code and save a lot of CPU time and memory – which are the primary resources for HANA.

Upfront, it was not possible to know that this would be the solution or by just looking at some performance trace. The primary developers had looked at this issue for months and could not figure it out. Instead, it was necessary to deeply understand what each part of the process is meant for and how it contributes to the performance. Performance optimization is just one part of my job. I also help development teams by picking the best way to implement a certain functionality on their data platform. Often, frontend- and API-/business logic developers don’t have a good overview of what tools a platform like HANA provides, and even less often they can evaluate which tool would be the best choice for a given task.

That is not the fault of those developers. After all, how can one single person stay on top of the whole development stack, where every layer has X components, and each of those components gets bi-monthly major updates? That’s where I help with the data processing layer. So, I’m probably not the best person to talk about the frontend framework, but I’d say I cover HANA and databases rather well.

2. Continuing, please tell me which departments do you work with most often?

Most often it’s analytics and data science teams that have their data in one or many databases and want to implement their solutions with it. Usually, this starts fairly easy, but once the reports, dashboards, and insights become important to the business parts of the solution may not hold up too well. Sometimes “big” queries use up all memory and render the HANA instance blocked for others. Sometimes the daily data loads do not fit in the allocated window anymore, etc.

This usually is because the analytics teams are focused on delivering the right semantics with their products and basically they have to trust the platform that it will deliver the required performance. There really is no good tooling for analytics and data science users to learn how their code/models perform and how to scale it easily. All that is available are DB/HANA expert tools.

My role is knowing both sides of this chasm and enabling a feedback loop that eventually leads to a better solution. And this is precisely what I offer with my company Data Process Insights. Getting analytics right, both on the semantic and the technical level, so that users, in fact, can learn useful and meaningful things from their data and make data-driven decisions.

3. That sounds interesting, especially concerning my next question. You write about yourself that you have extensive experience in solving difficult problems. So, what challenges do you face at work?

Okay, so one person’s “difficult” is the other person’s “bread-and-butter”, but if the core developers and lead architects of big data products struggle to make their product scale, then that is the kind of problem I tend to work on. Then again, the challenges could be not-technical in nature.



FlashPipe
Simplifying CI/CD
for SAP Integration Suite



SIGN UP NOW!

Smart Data Integration

Flowgraph data processing

INT4



Uladzimir Sakalenka
SAP Technical Consultant

[int4](#)

Smart Data Integration (SDI) flowgraph data processing

Smart Data Integration

source data preparation

INT4



Uladzimir Sakalenka
SAP Technical Consultant

[int4](#)

Smart Data Integration (SDI) – source data preparation



Automated testing

SAP PO

SAP CPI

SAP® Certified
Powered by SAP NetWeaver®

SAP® Certified
Integration with SAP S/4HANA®

EXPLORE NOW

SAP EIM - SDI



Health platform was a clinical data warehouse for millions of patients' data and thus very sensitive in terms of data security. As the platform was not designed as a SaaS solution, such a test drive would have required their own HANA server, that is too expensive, or cloud VM which is not acceptable for data security reasons.

The solution I came up with was to create the HANA-in-a-box instance that had the whole solution in a pre-configured VM running on a local Intel NUC machine. This guaranteed minimal deployment times and maximum control over the data by the clients. They could literally take the HANA system home and physically put it into a safe if required.

4. But you used to work as well for SAP for over 15 years, going through different positions in different countries. I'm wondering which of those countries do you remember most, and of course, why?

I started to work for SAP in Vienna, Austria, which is a great city. I met many good friends there and am looking forward to the time when we can travel again. Because of COVID, there's no travelling from Melbourne and Australia right now.

After living there for about 12 years, I moved back to my hometown which is Wuppertal in Germany since my work at that time focused mainly on international clients, and so I used that as an option to spend more time with my parents and old school friends. That was great, too.

And then, I met my partner and moved in with her straight to Melbourne. Well, Melbourne and Vienna share the same fame in having "won" the title of "most liveable city in the world" a couple of times, so I cannot really pick a favourite here. However, where we live now, I just have to cross the road to get to the sandy beach – that's hard to beat.

5. We already know that you used to travel frequently. Still, it is also worth mentioning that you are a co-author of the book "SAP HANA Administration", and a prolific blogger as well. Why are you so willing to share your knowledge and when will we see your new book?

Writing and publishing the book quite simply was a life-goal for me, and I guess my co-Author Dr Richard Bremer had a similar view. When we wrote the book in 2013/14 the state of SAP HANA documentation was not great, and there was no proper reference book available that we could refer our customers to. What we wanted was a book that would explain core concepts on a level that makes sense to the DBA or HANA user, and that would enable them to use it effectively. We got lucky by getting a fantastic editor from SAP Press, a big shout out to Kelly Grace Weaver, and the book fared really well. More than six years later, it's still being bought, which is astounding to me. It's such a niche topic, and the product moved on a lot since then.

A new book is currently not in planning. We have been asked a couple of times about writing a new edition. Personally, I would rather write a different book altogether instead. Writing my blog or answering questions is dear to me. When I started working with databases, the documentation was still in the form of thick paper-reference binders. It was hard to get answers to problems quickly. Years



int4

SAP EIM – Smart Data Integration (SDI) roles and privileges

Smart Data Integration

initial setup

INT4

Uladzimir Sakalenka
SAP Technical Consultant



int4

Smart Data Integration (SDI) initial setup



Then, working for SAP, I joined SDN, an ancestor to SAP Community, and started to answer questions myself. After a while, I wrote a first blog post, and everything else followed from that.

With my [blog posts](#), the main reason for writing usually is that I want to work through a topic for myself and not lose it afterwards. So I write it up and publish it. Thanks to google, I can now search for my old articles if I fail to recall specifics.

With answering questions, mostly on SCN and Stack Overflow, my main benefit is that many of those questions do not have obvious answers or present problems that are not typical for SAP environments. Getting “strange” questions is an excellent way to broaden the horizon and to check your own assumptions and knowledge.

6. Since you are on friendly terms with “HANA topics”, I’d like to know why SAP HANA is unique from your point of view?

The fundamental idea behind SAP [HANA](#) is to pack the majority of data management capabilities, an SAP customer from any industry would need into a single platform. Running the transactional ERP system on the same platform as the operational reporting, the data science-driven customer analytics as well as geospatial computing, and machine learning.

This “all-in-one” approach takes away many challenges that come with operating separate systems away and allows users to use another feature when they need it. For example, if you want to use special functions to process hierarchies or network data, then there is no need to install another product for that.

7. As far as Machine Learning is concerned, I assume that it is not a new concept for you. How do you think this area can be used when it comes to ERP systems?

The extent of what the majority of users do with their computers is frustratingly little. As Alan Kay demonstrated at OOPSLA in 1997 – that’s 23 years ago, longer than my whole professional career – “[The computer revolution hasn’t happened yet](#)”. While everybody in the well-off countries now has access to a global knowledge and computing network, the ability to use this and to learn and discover new things by using it has not progressed.

Much of organisational computing is still done on a simplistic Excel level. You know, sort

and filter data, maybe a VLOOKUP or pivot table for the power users and that’s about it.

I guess that this is partly because better methods are still too technical for the non-scientific, non-coding users and partly because there is a steep learning curve. Finally learning per se is an investment, and somebody has to pay for that. Now, as an organisation, I can either go and try to train my workforce in general methods, and the special applications that are relevant to their respective role. Or... I can teach a machine to do a specific job automatically.

that process has been done in the past without writing much code that needs to be managed and updated. If successfully applied, such solutions automate such jobs and the person that used to do it, is now free to focus on the tasks that require more knowledge, experience, and decision making than the “dumb” part that was automated. All in all, the Machine Learning we’ve seen in the past decade or so, is about the automation of pattern recognition – basically, it’s programming by example instead of code.

8. Okay, so the question that arises to me is, what is learning SAP machines? And for which industries can it be of use?

SAP’s offering in the Machine Learning space has been a mixed bag in the past.

Some parts of it are built into the SAP applications, like SAP Concur, SAP Hybris, or the various planning solutions. Other parts are pure technology offerings, like the HANA ML APIs/Libraries.

The former solves a specific function for SAP users. At the same time, the latter opens up the SAP data stored in HANA to a broader set of technology commonly used by data scientists and data analysts.

9. Last but not least, On-premise or Cloud? Which of these solutions would you choose?

To me, the answer to this question is mostly about managing costs for a solution. How much and when will I have to pay for a given solution or part of that solution. And also what quality of service can I get out of this?

Running your own IT infrastructure can easily become a person-intensive and upfront cost heavy task that, in itself, rarely helps to make money. To many organisations, IT is still seen as a cost centre and associated with “taking care of the servers”.

Beating the costs and quality of service of any of the cloud providers is hard and probably not something most companies should want to invest into – that’s why in my mind using a cloud-based operation is the default choice when the organisation has at least some experience with it.

With SAP customers, the situation is somewhat special due to the fact that many of them have already done massive on-premise investments, both in hardware and staff, which makes those investments “sticky”.

A good analogy probably is to see how industrial power changed from something that would be generated on-site in or near the factory, and that is now routinely simply consumed through the public power grids.

This is the exact idea that John McCarthy explained in 1961, I mean, computing as a utility. It’s kind of disappointing to see that the “big ideas” of today really are mostly very late realisations of the ideas from half a century ago – especially for an industry that claims to be the innovation motor.

Read also:

1. Autonomous system such as Tesla autopilot could be a future of the ERP
2. SAP Application Server is like Arnold Schwarzenegger in the Commando movie



Share:   



Jarosław Ziółkowski
Sales Executive

Sales Executive who has over six years of experience working with Enterprise Customers in the IT industry. He would like to do thousands of things at once, but he hasn't yet found a way to extend a day. An analytical mind, focusing on the most exquisite details. Although he loves to talk a lot, he listens much better. A journalist by education who daily exploration of the secrets of SAP integration





+41 41 562 02 21

HOME

SOLUTIONS

INT4 IFTT

POSTING TEMPLATES

S4MIG

CONSULTING

COMPANY

ABOUT

CAREERS

LEADERSHIP

BOOKS

PARTNERSHIP

BLOG

VIDEOS

CONTACT

INT4 IFTT TRIAL

2021 copyrights INT4 | [All Rights Reserved](#) / [Privacy Policy](#)

