



For Cloud Professionals, part of the On Cloud Podcast

David Linthicum, Managing Director, Chief Cloud Strategy Officer, Deloitte Consulting LLP

Title: **Google and Hadoop: helping banks meet their challenges and build a brighter future**
Description: Large financial services institutions (FSI) have traditionally been reluctant to fully embrace the cloud. However, as data volumes expand, customer demands grow more complex, and regulatory requirements increase, many FSIs are finding that they have no choice. For many FSIs, Google Cloud's new tools and technologies, coupled with Hadoop, have made cloud computing more attractive and mitigated barriers to entry, but challenges still remain—especially with DevOps and talent. In this episode, David sits down with Deloitte's Sandra Bauer and Google's Apurva Desai to discuss banking and the cloud and how Google's cloud offerings are helping banks leverage technologies like AI and techniques like DevOps to help build a better, more scalable, more secure future.
Duration: **00:20:45**

David Linthicum:

So, welcome back to the podcast. Let's talk about cloud computing and banking. I'll tell you what, in all the cloud solutions I built in my career, half of them have been around banks and bank computing, and so this is kind of a topic that's dear to me. Joining me we have our special guests. We have Sandra Bauer with Deloitte and we have Apurva Desai with Google, and we're going to talk all about banking. But let's first learn about you. Apurva, tell us what you do at Google, and how you got involved in cloud and banking.

Apurva Desai:

Sounds good. So, I am one of the engineering directors and I manage open source data analytics, so products like Dataproc, which is a managed Hadoop and Spark offering on Google Cloud; Cloud Composer, which is our managed airflow offering; and we are actually announcing a beta release of an acquisition we made last year with Cask technology. It's a managed CDAP product called Data Fusion. I've been with Google Cloud for about three-and-a-half years, and as a part of – but I have been involved with the Hadoop ecosystem since 2009, when I was with Yahoo! and

Hadoop was in its infancy. So, I've seen it kind of grow. As a part of my experience with Google as well as with prior companies, I have a fair number of experiences dealing with enterprises, including banks. So, this is an exciting field.

David Linthicum:

So, what attracted you to working with Google Cloud Platform?

Apurva Desai:

So, when I actually was looking around, I was being fairly selective about the opportunities that I wanted to perhaps pursue. Google Cloud about three-and-a-half years ago was fairly up-and-coming. So, this looked like a greenfield opportunity. Cloud was growing. The charter had just about finalized. The structure was being put together. And I got to grow some of these products pretty much from their infancy, all the way to where it is right now. So, a very good time, good place. I guess right time, right place.

David Linthicum:

Sandra, tell us what you do within the firm and how you came to get introduced to Google Cloud Platform and working with banking systems, things like that.

Sandra Bauer:

Yeah, sure. I'm a director with Deloitte in Germany, actually. Deloitte Consulting in Germany. I've been working in the field of big data and analytics and AI for the last three years. We were implementing data platforms for two of my biggest clients in Frankfurt, two big banks. And I got into the, or starting, the Google Cloud Platform Alliance in Germany, taking it over from the UK, and doing the training of our practitioners myself. And my biggest client right now just signed, two weeks ago, the Google Cloud contract, and we can scale out their big data platform into the cloud now.

David Linthicum:

Outstanding, outstanding. So, what are the typical use cases you're running into? I'll start with you first. What are you seeing?

Sandra Bauer:

Well, from a banking perspective, and I'm doing FSI as a specialization right now for the last three years, it's mostly around regulatory and compliance. So, a lot of banks are putting data lakes in place to kind of push all the data that they have to report to the authority through the data lakes and get that kind of stabilized and use it as a data hub.

Going forward now, what we see a lot is AI use cases coming up. So, not so much the data engineering parts, because the data sources are pretty much streamlined through the organization already, but on top of the data that exists, they want to build new models, to better serve the customer, for example, for Customer-360 purposes, to upsell, or cost-reduction use cases. And we also have a category where they actually do innovations.

David Linthicum:

So, it would be compliance typically. It would be regulatory reporting, which I guess is compliance. It would be auditing, internal systems, trading systems, things like that. Where is the bulk of the work going right now?

Sandra Bauer:

Right now, we are looking pretty much into cost reduction, because the banking industry is heavily under constraints in budget purposes, but also increasing the revenue share.

David Linthicum:

Cost reduction of IT?

Sandra Bauer:

Cost reduction of IT, or the back-end systems. Like, a lot of robotics use cases are coming up and streamlining the back-ends and the customer interaction things.

David Linthicum:

Back office processes.

Sandra Bauer:

Back office processes. Chat bots are coming up dramatically, actually in all the areas, from customer interactions with a customer service desk, or in the HR area for the internal employees. But, also, looking into the more profit-oriented ones like upselling, cross-selling, doing different services that are not specific to the FSI industry, because FSI is struggling, right? What do you need a bank for? My transactions, I have personally only used the bank over the last year for getting cash. I'm not stepping into any branch anymore. I'm not interacting. I do everything online. So, with these systems coming up, the banks really struggle with their core service that they actually do for customers.

David Linthicum:

Those are some big problems to solve. Apurva, why is Google Cloud Platform the place to solve those problems?

Apurva Desai:

I think perhaps to build up on some of those use cases, one of the interesting ones we have seen of late is around fraud detection. This is where Google Cloud actually can provide a lot of value, because our streaming portfolio or our ability to do near real-time, even if you don't go fully real-time, is quite outstanding. We have a lot of products that provide that. Specifically, in the context of Hadoop, we provide – or open source, we provide Spark Streaming.

Just to give a clear example, oftentimes you are seeing now when you do a transaction that looks somewhat iffy, the bank would send you a text saying, "Is this right? Is this wrong? Did you actually make this purchase?" It makes no sense to send that a month later, or a week later, or a day later, saying, "Did you make this call?" You want to make it as real-time as possible.

And this is where Google Cloud actually comes in very handy. Our managed offering called Dataproc, which is a fully managed Spark offering on top of Google Cloud. You can use Spark Streaming to pull in a lot of these sources of streaming data, make some simple models or complex models using that, and put out the output back into the source, saying, "This actually has some form of issue."

The streaming, or near-real-time streaming, is something which is becoming very quickly a differentiator for Google Cloud. We have a fully, fully managed serverless offering called Dataflow, which also plays a pretty big role depending on how the banking system wants to graduate it.

David Linthicum:

So, do we have a scaling limit to deal with? What's a typical scaling limit in dealing with Google Cloud Platform? Apurva, I'll go to you first, and then I'll let you follow up, Sandra.

Apurva Desai:

That's an interesting question. If you start breaking down the problem, you're looking everywhere from network, down to storage, down to compute, down to analytics. These are the four primary components to come together to building an ecosystem. Some of these I think are fully managed and fully controlled: compute, storage. Network oftentimes gets limited by what a customer has on their source end. So, we're only restricted by whatever they can pump in.

The integrated solution oftentimes runs into some form of a scaling issue that needs debugging, but individually I think Google Cloud provides probably one of the most unprecedented scales, because our offerings are pretty much tested internally at Google's scale before we actually push it out. A lot of the cloud-native products we build have been tested at Google's scale internally, so we feel pretty confident about that part.

Open source is where we play nicely with the open source ecosystem. So, whatever are the limitations, we try and make it – mitigate it by integrating them with a cloud-native solution, but there are obviously some limitations that we just cannot overcome.

David Linthicum:

Sandra?

Sandra Bauer:

One of the biggest challenges that my clients have right now is their on-premise Hadoop is not at the scale that they need it for certain jobs during the month. So, one of their jobs that's running once a month actually –

David Linthicum:

Are they running out of capacity, computes?

Sandra Bauer:

They're running out of compute capacity, yeah.

David Linthicum:

Compute capacity.

Sandra Bauer:

Compute capacity.

David Linthicum:

Not storage capacity?

Sandra Bauer:

Not storage, no. It's just compute, and it's just once a month, right? Why would you size your system for a job that runs once a month? So, what they're looking into is upscaling the systems and going hybrid between on-prem and then into the cloud with those compute nodes.

David Linthicum:

So, the existing as-is stuff and Hadoop is on-premise?

Sandra Bauer:

Correct.

David Linthicum:

Got it.

Sandra Bauer:

And for those kind of things – and they're coming up much more the more you actually put data on your Hadoop system. People come up with all different sorts of utilizing that data and multiplying it, and then getting new data sources onto the Hadoop system which they just use once, for example, to prove something or to do an analysis on social media, for example, and then discard it again.

So, you don't need to keep that dataset because it's not worth afterwards anymore. For example, if you have a big event like a speech, for example, where somebody wants to see the web traffic – what the speech actually produces as web traffic, once the speech is done, yeah, sure, you can keep it for reference, but it doesn't help you. So, those computes need to scale up and down fairly quickly.

David Linthicum:

So, I guess you're playing Monday morning quarterback on this. It's an American term. So, what went wrong? Was there a lack of planning? Was the technology the issue, the database model?

Sandra Bauer:

No, it wasn't the technology, and it was not the planning. We were fully aware that we didn't scale it to the possibility because they do what they think they need for regulatory purposes, for example. But everything on top is basically nice to have.

So, they were fully aware when they put those data platforms in place that they were shortly running out of storage and out of compute, too. And then they thought, "Well, once we did the first step and we sell that to the customers, the demand will grow. And once we have the demand, we can actually put more budget requests in and then scale up."

And, also, at that point when we started three years back, they were like, "We're never going to go to the cloud." None of our banks that I'm working with was really seriously looking into cloud, except for software-as-a-service. But for platform-as-a-service they were like, "It's not secure enough. It's not in our datacenter. We don't have the controls over the keys. Everything needs to be anonymized. And then if somebody gets access to it, we don't want to allow that." But now three years later they see that the capacity is just not there, and not scaling into the cloud is not an option anymore.

David Linthicum:

Got it. So, in other words, they really kind of have no choice. Ultimately, they have to look at the cloud technology to save them and can do so typically at a more economical advantage than it can ever just moving forward with some of the on-premise systems. So, what would your advice be to a client like that that needs to scale out and is looking for your technology to help them?

Apurva Desai:

I think one of the things that we constantly see with clients who have started on-prem are the issues with DevOps and DevOps-related challenges. I think Sandra referred to it a little bit, but everything from capacity planning. How much do I need? Do I have provision for my peak capacity? Do I have provision for my average capacity? What about my network bandwidth? How do I procure these machines – cable, rack, stack – provide accounts, put the right operating system on it, put the right patches on it, create the necessary queues so that people can do ad hoc queries?

All of those issues are more or less taken care of for you as soon as you come onto the cloud. So, the journey to cloud oftentimes becomes so much more streamlined for the DevOps team that even though they are the first teams that somewhat have a resistance, they're the first teams that actually embrace it, saying, "This is actually making my life so much easier because I don't have to do a 12-month capacity planning or a CAPEX planning. I don't even have to do any of the upfront provisioning. It's on-demand: when you need it, you have it."

There are a few other things that you get for free along the way, because we have on Google Cloud something called a universal authentication authorization system. You get the audit logging and the SOX compliance-related work, somewhat for free because we know who did what at what time in the system, and you can actually get a pretty good stack trace of that. So, these are some of the things that make migration or journey to Google Cloud that much easier and compelling for the customer.

David Linthicum:

So, we've got this concept called DevOps, and every time I'm doing a cloud project, I'm doing a DevOps project at the same time because we're putting in advanced cloud technology such as Google Cloud Platform, and the ability to leverage that effectively is going to be dependent on my ability to change and build software at the speed of need, that those businesses, as they need it.

It seems like banks, certainly in the United States, they have to change all the time, getting new factors, regulatory compliance, new products out and running, and also really trying to keep up with the competition. Really, cloud computing becomes a force multiplier in the market.

So, what are the dynamics of looking at DevOps and cloud computing, and how is that working in the banking industry? Sandra, I guess I'll go to you first.

Sandra Bauer:

Sure. We actually see a really big push towards DevOps in the projects that I'm doing, which is also a very big challenge for the clients because the people that they have are, first of all, most of the time knowledgeable in their old environments, but not so much into the cloud.

Being agile, and developing in an agile way, doesn't come natural to them. They are used to their lists of things that they need to work through and then somebody telling them what to do, and also the flexibility that you have in deciding by team what to do first without actually neglecting who else is depending on you, that needs a lot of orchestration. That right skillset and the right people to actually push that through your organization is one of the biggest challenges that we have.

David Linthicum:

So, any special needs with DevOps and Hadoop and the ability – and even leveraging the cloud? So, you create a tool chain really with the target database and target cloud platform in mind. Is that what we do?

Sandra Bauer:

Well, basically you look through what kind of services you want to provide, right? And then you build the whole value chain along from extraction into the data modeling part and then the analytics on top of it. And then you do that hopefully with the same team over and over again, so you can become faster in developing things. But at the end of the day, you also need to have the teams in the business who understand what you want to do, and how they should be using it so that the upscaling of not only the DevOps people, but also your whole organization becomes a very critical part going forward.

David Linthicum:

So, I'm finding DevOps, which is development, as a means of operation and the ability to really take down the barriers to do it and work with agile, and the automation of agile, is a difficult thing to implement within lots of enterprises, specifically banks, because it requires a cultural change. That's probably the most important thing to do. You know, building the DevOps tool chain is fairly easy.

So how are you guys facing the cultural problems in these different organizations who are looking to use Google Cloud Platform? How are you getting around that obstacle?

Apurva Desai:

I think you actually nailed it. Oftentimes you will hear this term, "lift and shift," and oftentimes people equate that with lift and shift of technology. For me, it's actually way more than that. Technology is perhaps the most mitigatable, or manageable, piece of the tool chain. It's about people. How do you actually get the right mindset going? And, in this specific context, if you talk about DevOps very specifically, the mindset changes from day-to-day operations to looking at operational metrics. This is a journey. It doesn't happen overnight. The way a lot of the times the customers end up doing it is to start seeing incremental value.

If I do not have to build my infrastructure, and if I am getting a burst capacity pretty much automatically on Google Cloud, the way I monitor and operationalize my operations, I guess, is very different on Google Cloud because I get Stackdriver logs or I get built-in logs. I get built-in ways of figuring out who did what, when. I have a really clean way of augmenting my capacity.

So, I think the mindset changes to saying, "Let me build alerting and dashboarding on top of the data I already have, versus data I have to actually mine and build." That journey actually, sometimes it takes a month. Sometimes it takes a little bit longer.

But as a part of Google Cloud, our professional services organization, our customer engineering organization, our PMs, the engineering org, which is what I represent, oftentimes get with the customers to say, "Let's sit down. Let me address the fundamental issue is your job isn't going away. It's just taking a different dimension. Let me help you map that." And once people see it, I think the journey becomes much more palatable.

David Linthicum:

So, is there often a change in personnel, or a change in training plans, or skillset divides that are conquered in doing that?

Apurva Desai:

All of the above. Oftentimes what we have seen is a little bit more of training and getting them over the hump of, "Am I losing my job?" does the job a lot more effectively than – and most people want to learn. I haven't come across many orgs where they're like, "Oh, this is new stuff I don't want to learn."

So, if you provide them that learning opportunity, we have a lot of Google classes that we offer, a lot of training on-prem that our training personnel can go and certify people. Once they see that, I think we have had little resistance. Then the question changes to what I mentioned earlier, which is a much more manageable issue, of just technology migration. That we all love, because that's a cookie-cutter approach we can apply for a lot of those things.

David Linthicum:

Sandra, if someone's going to walk up to you at Google Next, and they're a big international bank and they have a few SaaS systems. Most of their stuff is sitting on-premise. They figure their IT is about 50 percent over comparable cost within other companies, and they're trying to reduce that down.

So, what would be the first three things you would tell them to do to start moving in the right direction, besides hiring us?

Sandra Bauer:

Well, of course hiring us, and we do, what we say, a cloud assessment with them. We have a whole catalog of different varieties of things that we have to look at. So, you look at the strategy for cloud. You look at the people, the skills, the business processes, and the technology, of course. And then based on that assessment, we actually give you a roadmap on where to move first, where the quick wins are.

We also put a training path into the roadmap to upskill your people, or to hire new people for the skillsets that you need. And then just a basic roadmap on where you want to go as an enterprise, right? Do you really want to do this, or do you have to re-emphasize your business model and maybe do something different, some new services that you offer your clients?

David Linthicum:

So, what would be the kind of mistakes that are being made that enterprises are making? Anybody can respond to that one.

Sandra Bauer:

Well, usually, and I'm from Germany, what we see that usually helps, talk to Google, for example, talk to Deloitte, and ask them about our experience that we have. And we also of course do consulting work, but at the end of the day, it's about the relationships. We want to make our customers successful.

And the most successful project that I have seen was actually when we work together with a vendor like Google who has the same mindset as Deloitte and Google and who really want to push things, too, and do new, innovative, and exciting things. Just talk to us and ask us, and we're more than happy to help.

David Linthicum:

That's great advice. We're going to go ahead and end it there. I want to thank my guests, Sandra and Apurva. I appreciate you very much for showing up on the podcast. Some great advice on how to leverage cloud computing for banks. Thank you very much for attending the podcast. You guys take good care.

Visit the On Cloud library

www.deloitte.com/us/cloud-podcast

About Deloitte

As used in this podcast, "Deloitte" means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Please see www.deloitte.com/about to learn more about our global network of member firms. Copyright © 2019 Deloitte Development LLC. All rights reserved.