



The Deloitte On Cloud Podcast

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

Title: How leveraging data can help you enhance cloud performance and value

Description: Though cloud adoption is rising, some enterprises aren't seeing their desired cloud value and performance - but the problems aren't insurmountable. In this episode, David Linthicum talks with Snowflake's Jeff Hollan about how companies can improve their cloud operations and gain more ROI from cloud by leveraging their data effectively and utilizing a data cloud. Improved data strategy and governance, low data latency, collaboration, and giving developers the support they need are also imperative.

Duration: 00:26:10

David Linthicum:

Welcome back to the On Cloud podcast. Today on the show I am joined by Jeff Hollan. He's Director of Product at Snowflake, leading the developer platform team. Jeff, welcome to the show.

Jeff Hollan:

It's fantastic to be here. Thanks much for having me.

David Linthicum:

Hey, Jeff, where you calling from? Where are you physically located?

Jeff Hollan:

So, I am a little bit outside of Seattle here at the Snowflake office, and it's actually a beautiful sunny day. Usually, we're in a bunch of darkness and rain right now, a lot of clouds, but it's gorgeous outside.

David Linthicum:

Yeah, it is in Virginia too. It's 55 degrees. May go out and even take the bike out. Can't believe that. So, it's great to get into spring as soon as we can.

So, give us the Jeff story. How did you get to Snowflake and where did you come from and what'd you focus on your career? Love to hear that.

Jeff Hollan:

Yeah, sure. So, I've always been really passionate interest in technology. Going through university, I remember at the time my part-time job while I was going to classes was managing some servers and workstations for some small companies, which was fun. And pretty shortly after university, joined Microsoft, and so that's where I spent about the first decade of my career. And right around, I don't know, a few years into Microsoft is when I moved into the Azure cloud side. As a developer, I'd been hearing so much about the cloud. This would have been, I don't know, fifteen years ago, ten years ago at this point. And it moved in right at the time when, on a product standpoint, Azure was starting to move into serverless. I don't even know if serverless had been fully defined in that world, but those were the first products that I was working on, so like the serverless workflow solution, later on to go into Azure functions and drive all of Microsoft data as your serverless portfolio.

So, that was a ton of fun. I really loved working in that space. Took me into a bunch of different corners including understandably, when it comes to serverless, things like containers and cloud native through some of our explorations and investments had created in helping open source a few pieces of technology that are in the cloud native compute foundation, and just all things developers. I was a huge fan of finding how can we unlock the value of the cloud for any developer and let them bring those best practices and that optimal developer flow into building cloud solutions and serverless, and containers played a huge role there. So, that's what I was doing beforehand.

And what's interesting, actually, when I was thinking about chatting with you today, I was remembering in the last few years while I was at Microsoft, we were noticing a few trends. One of them was that what, traditionally, I would describe is somebody would come to the cloud, and they would say something like, "I want to build a website." And, so, they would start building the website, and then at some point down the line, they're like, "hey, I need a database to store all of the customers who are coming to my website." So, then they would go figure out what flavor of SQL or whatever else do I want to have.

I was seeing an increasingly large set of applications who were coming to serverless, and it was organizations who were saying we have all of this data. We'd been collecting all this data from IOT devices, app telemetry, you name it. We want to build serverless solutions that are helping make sense of that data. We actually want to turn that into insight and value. And, so, it almost flipped. There was this data-first application-second mindset rather than application data first, and this was something that we'd been talking about for a while.

In the middle of all this discussion, someone reached out to me from Snowflake and was like, "Hey, Snowflake's been around for a little bit, they're in the cloud space as well, would you be interested in moving to Snowflake?" And at the time, I was like I love what Snowflake is doing. I equate Snowflake very much with data, and I was like as cool as that is, I'm also super passionate about these developers, like these VS code, these Python, you name it, developers. And they're like, no, that's actually great because what we're seeing is the same thing that you're seeing and that people have this data and now they're trying to figure out how to make more sense of it. So, early last year, I made the move and I joined Snowflake, as you mentioned, leading their developer platform and compute services teams, helping it so that folks within Snowflake now can have some of those same serverless benefits in building applications.

David Linthicum:

Yeah, I love what kind of motivated you there because I'd see the same thing. People weren't weaponizing data, not leveraging it as a true force multiplier for the business, and now that we have AI where the data becomes the single most way to train these systems to do these amazing, magical things, And, so, the data becomes the center of the universe when people are looking at how to improve their existing IT systems.

David Linthicum:

So, Jeff, one of the questions I think that comes to mind that I think a lot of the listeners have right now is, what is a data cloud and generally what is it used for?

Jeff Hollan:

Yeah, so data cloud, the way that we like to think of it is it's a cloud first way to build cloud applications that is really centered around data. So, I guess in some ways it's self-descriptive, but often when folks start thinking about the cloud, there's the big, big clouds –the traditional clouds – very infrastructure focused, provide those virtual machines and some PaaS services, and then there was a flavor of this even when Snowflake started where folks were like, okay, so Snowflake is like data warehousing in the cloud. Is that the way that I should think about it? And the reason we say data cloud is because we want to be very clear. No, it's so much more than just doing some of those traditional warehouses or analytics workloads.

It's really about providing cloud functionality and cloud differentiation and cloud innovation that, yes, is center and core and centralized around data, but a broad set of capabilities from building visualizations to building machine learning models to sharing business insight or even running aspects of line of business applications. So, that's why we like to talk a lot about data cloud it's not just some of those traditional, "Oh okay, I have my product data, I need it to show up on my dashboard," and that's the only thing this thing does. It can do so much more than that.

David Linthicum:

So, using data kind of as a multidimensional concept. In other words, it's not just persisting data for analytics, decision support, things like traditional data warehouses do, but also supporting transactional systems as related to business intelligence and really kind of taking information, putting it in a single place, and having multiple ways of leveraging that to add value to the business. Is that close?

Jeff Hollan:

That's spot on. And it's really, I think, a side effect of the fact that data has always been important, it's always been the lifeblood of every organization. Increasingly moving forward, especially as we look ahead to AI and ML, data is just so critical to everything, and so it starts to weave into so many different workloads, and so it's super useful to have a cloud that doesn't let you, doesn't make you silo in those data pieces, but it can weave that data throughout your applications and it's much more expandable.

And just looking ahead to, as you mentioned, all of these major trends that are happening, data science, AI, data gravity, cloud native applications, and saying okay, data plays a core role into so many elements of an organization and to a business, but Snowflake isn't just going to be the query engine, the analytics engine. It's really looking to say, "Okay, how can we branch out into providing more capabilities to say run code logic, to do that application sharing as well," so it's just a much broader definition, and it starts to make sense as well, as you mentioned, when you look at trends like AI, saying, "Okay, yeah, data plays this fundamental role in being differentiated," but it's not just going to be about data. It's what can you do with that data. So, Snowflake's evaluating and working with customers to figure out "okay, where is that going to take us, where should we go next."

David Linthicum:

Yeah, it's playing 4D chess with your data at the end of the day, the ability to figure out abstractions and automations and the ability to find intelligence and to leverage that intelligence into additional intelligence and building the knowledge models. People don't really getting that yet. I think there's a few companies that are innovative that I'm seeing as starting to exploit that, but you think about it, we're going to have to, as a group of businesses, provide a better customer experience, provide the ability to have insights into what our data means, the ability to drive smart supply chains, all these things that we need to do better as enterprises, as businesses, and that's really kind of foundational in the data.

So, what considerations should people who are building these things, cloud practitioners, take when they want to create an ecosystem unique to them to build an application?

Jeff Hollan:

I guess the first one, and anyone who's been in the space will be like we already know these things, but it's worth repeating again. Making sure your data is quality data is a huge, huge consideration, and so you can't take your data for granted just because you're collecting data. You have to figure out how to make sure that you've got the right controls, that you're testing the validity of the data, and keeping track of it. Like anything in technology, there's very rarely a set it and forget it. You're going to want to make sure you're having a plan and the right technologies or partners so that you can keep an eye to make sure that your data's not drifting, that the quality is up to par, that you have the right processes and procedures in place. So, that's one big investment for practitioners to consider.

And the other one I would mention as well is figuring out – again, I think that there's so many parallels happening in the cloud world, and there's another parallel that is an important consideration to have is around DevOps. And the reason I say DevOps is, at least for me, one of the things I've seen in both the development world and the cloud world and seeing now in the data world is this really big need for organizations to need to move faster from here's the need that we have, here is the gap that we have, how can we get that solution out into the hands of our users and to our business as fast as possible.

A lot of that is going to come down to the tools and processors and procedures that you have to go across that journey, to go from development to production, and so another huge one is don't think that data is any different. In the same way that when you're building containerized applications, that you're hopefully thoughtful around what's my CI/CD process, what's my testing process, what's my deployment process, the same rules hold true for data. How am I ensuring that I can rapidly iterate on my data, move those changes through safely into a production environment, and then evolve and adapt? So, those are the two, I think, very top of mind that I would have. Data cleanliness and quality is something you want to keep track of, but then having the right processes to help you be productive so that you can really see those benefits from data in whatever cloud technology you're using.

I think one of the challenges that every organization is dealing with right now as they think about their data strategy is that how can you better leverage data to provide real value and ideally competitive differentiation in the business. And I think one best practice here which I would want to call out is it's important to have some sort of a data governance or data operability strategy in mind, trying to figure out how do you make sense of all of the different various signals that might be coming in. One of the things that the Snowflake CEO, Frank Sloatman, likes to say is, try to keep your data from landing in a silo. There's other aspects as we've been moving forward to the cloud that might have been siloed or siphoned away. Your data cannot be in a silo within an organization, and that requires some forethought. And it's not an easy problem to solve because that volume of data is just growing exponentially. So, that's one thing that I always encourage, and I think we've seen organizations who are successful in their data strategies is they are very thoughtful. They put data and data strategy as something they think about ahead of time rather than just being reactive and trying to just herd a bunch of cattle that have already run away.

The other aspect that I would talk about here is 'figuring out how you can bring some of your application logic as close to the data as possible. This is across cloud providers a number of different trends. One of the ways you'll see this being talked about is okay, when you have things like IOT devices, you might be thinking about bringing your processing to IOT devices, and that's a real trend.

But even once it's in the cloud, once your data is there, there is real cost in latency that comes from moving around your data too much, and so that's really what's motivated a lot of Snowflake's development platform is these organizations who just have these massive mounds of data and they're trying to figure out okay, as we're thinking about maybe training a machine learning model or enriching our data or pushing it downstream to another partner, how can we do this without having to move the data and deal with the latency and the costs of moving those things.

David Linthicum:

Yeah, I love the line from your CEO that don't put your data in silos. It seems like we've been trying to knock those things down for the last 20 years and aren't having much luck, but the thing is when you're in a hole, stop digging. When your ability to kind of open up your information and make it more accessible and viable, and also I like the fact that you're building governance systems. Data governance is something that's often overlooked but it's a key chain of custody, ability to put policies around utilization and data, objects and tables, things like that is really kind of fundamental to when you get to a data architecture. So, what I heard when you were saying this, a lot of fundamental architectural best practices are starting to be practiced out there as people are building their cloud architectures. So, let's talk about data and developer components that enterprises are most likely to employ out there. What is your suggestion to people who are looking to move into this direction as to what are the major components they should be looking at?

Jeff Hollan:

So, here's something I'm really passionate about on this front, which is it's been fun coming from that cloud native Kubernetes container world. And there, just the overall DevOps life cycle or DevSecOps or whatever version of DevOps you want to talk about is very, very popular. What's been interesting when I talk to organizations who traditionally they've got teams who might be doing data analysis or data engineering or even data science, previously they would almost not follow those same DevOps practices. It's not uncommon at all to find an organization who the way they're building a mobile application, they've got a CI/CD pipeline and they have real-time telemetry and they have regression testing and all these different layers in that release process, but then when you jump to the other side of the organization and you talk to the data team, they may be like, yeah, we're just kind of making updates as we go, hopefully not deploying straight to production, but something close to that.

So, I think one component that is definitely growing now, and there's a number of partners and tools who are helping fill this gap, is how can you bring that same set of protections around DevOps to a broader set of workloads, even things like you've got an ingestion pipeline that's pulling data from a data source and making it available to your organization. How can you add DevOps to that process so that you can validate changes before they go to production so that you can have reviews and be able to roll back changes? The part that I chuckle on this is those types of conversations were fairly well underway in the serverless and container world, but they're earlier in the data world, and understandably so, but it's an important trend, and I think there's a huge opportunity moving forward for organizations to figure out how to make that more common.

And I guess this just comes to another important component, which is the lines are blurring across the board. The longer—I think one of the things that cloud has accelerated is that you no longer have just a single hat that you're wearing. You're rarely ever just a database admin or a front-end engineer or a back-end engineer or a data scientist. I think moving forward we're going to see data and AI and ML and DevOps span into a larger portion of roles, even into the citizen developer type role that you might think of how can that type of person also start to leverage some of these best practices? So, that's a big component that we are investing in and that I think you'll see a lot of other people in the space trying to figure out how you can add more of those best practices into more and more workloads and for a broader set of users.

David Linthicum:

Yeah, I'm seeing more interest in that. There's always talk about the citizen developer ever since the COBOL days. That's why COBOL was created. But the thing is ultimately providing simplistic tools that are able to abstract a lot of complexity from the capabilities of these end users to leverage their data in some sort of way where it allows them to do their job much better without depending on a database administrator and database developer or developer in general to go ahead and get the answers that they're looking for. And I think that's huge because ultimately people are going to seek out their own answers, and their ability to build their own systems.

And we have low-code systems and no-code systems and all these things that are coming along where it's truly going to empower people to kind of take their jobs to the next level by developing and innovating their own stuff, even though themselves are not programs and coders, but they do understand technically how their job works and therefore they can design their systems to best automate and innovate what they need.

Another thing related to practitioners is they want to know, what are some of the emerging best practices to think about when we start deploying this technology?

Jeff Hollan:

For sure. And it's evolving. That's what's exciting is the best practices, there's always new ideas. Data mesh is an exciting notion that has been spinning around in a lot of minds, which is – has a lot of parallels to microservices. How can you make it so that individual business domains can be responsible for their data? There's some best practices emerging there, so it's a fun space. It's definitely an exciting space for builders of any kind.

David Linthicum:

So, considering the life cycle of a business and they're just getting into technology and to the enterprises that have been around for 100 years that are on their 20th generations of technology, what is the best time to deploy a data cloud for an enterprise?

Jeff Hollan:

Now. Yesterday. I mean, generally, every organization's probably battling some of the same stuff which is there's just a lot of data. I mean, big data's not a new term in any degree, but now just the amount of data, the amount of data sources, the amount of data we can pull from the ecosystem is huge. So, that alone is a concern, but it quickly comes into the how do I govern this data. What do I do to make sure that different departments can own and collaborate on their data together? One of the important things too is it's not just about my organization but how can I become a part of a better data ecosystem.

So, I think fairly early in the process, whether you're even a company that has 20 individuals or 20,000 individuals, if these are the types of problems that you're having to start to think around it's like we want to make sure that our data's not in a silo, we want to make sure that we have a central way to govern this that is extensible and leverageable by different teams, by different partners, data cloud is the thing you want to start looking for because what you'll find you'll end up doing without a data cloud is that you might end up inadvertently just building these silos.

You build one-off solutions, you build one-off environments that then you now realize, oh shoot, we have this pocket of data here, we need that pocket of data over here, and now we have to do a huge lift to get those things talking to each other. So, as all things, it's something to consider earlier in the process than later, but there might be instances where for you maybe you're not there yet. I would just say consider it. At least be thoughtful. Make sure that you are aware of some of those big governance and data collaboration and data management things that every organization will face with, and if you're not going to tackle it right now, at least have it on your mind, on the radar to make sure you don't paint yourself in a corner.

David Linthicum:

Yeah, I agree with that. I think that data's important to every organization no matter what they are in their life cycle, small, medium, big, large. And the thing about this is also it allows smaller businesses to punch above their weight, the ability to kind of leverage data in such a way where you're truly going to be a force multiplier, an innovative differentiator for them to kind of take their business to the next level. As a former CTO in smaller product companies, things like that, I did look at ecosystems where I could be more valuable, and those are kind of the partnerships I wanted to create, and these of course in the days before the cloud. Now that we have the cloud, I can understand if people are putting lots of information into your systems, I'm going to be interested in building tools and technology to add value to that existing value. So, how does that typically work? How do you typically onboard a partner and a customer and ultimately is there a marketplace that they're in or something like that where they can be found?

Jeff Hollan:

Yeah, there is. So, we've got phenomenal teams. We've got a partner alliances team who does a great job. I've been reached out by a few upcoming startups who are wanting to invest heavily in the AI space. We figure out how we can make them successful and figure out where we can do things like co-marketing or shared goals or you name it, but I think the big thing as you alluded to is that idea of a marketplace.

What's interesting, and this is thinking a little bit more strategic as well, is I almost think about this ongoing conversation that's been happening around multi-cloud and data cloud and discoverability and data gravity, all these kind of things pushing toward this world where the importance of having a solution whereas a partner, I can build something that can be consumed whether or not the consumer is in any of the major clouds is super important. It reduces the amount of complexity that I have to write when that multi-cloud capability is checked. And as well as to the point even if a single organization might have chosen a single cloud, it's very possible that some of their partners or stakeholders might have chosen a different cloud.

So, this new class of organization and providers, and I know you've talked about on your show before, that sit one layer higher than the infrastructure cloud providers themselves, and the analogy I think of is with cell phones and with infrastructure. I remember a time when smartphones and who knows—I guess it was 2G at the time, I don't even know – but all the ad campaigns at the time were all talking about look at the map, look at how much coverage we have. We have all this coverage, we've got—here's all the spots you'll get cell phone coverage. And then on top of it, you get the coverage you wanted, then you would go and buy some device, and it would do what you needed it to. It would do your text, your calls but then there was this huge disruption that happened when smartphones came around, and suddenly you get a provider who says, you know what, we're going to build on top of all of this infrastructure, but we're going to offer this really curated experience, and to the point you were mentioning, this app store, this marketplace to actually exchange these very valuable capabilities.

And I think a similar thing is starting to happen in the cloud world right now, , there's been so much acceleration in the cloud but there's this layer on top of it, this smartphone layer that builds on top of that infrastructure but that no longer becomes your core concern. You can just start pulling in the capabilities that you want, and so that's why this Snowflake marketplace for both data and apps is such an important strategic opportunity for us, and I think for all organizations to think about where that would shift as multi-cloud and these higher order supercloud, metacloud, whatever word you want to use start to come into focus.

David Linthicum:

Yeah, we're taking the innovative nature of the development, the startups out there and even companies that are restarting things inside of big companies, and the ability to take this innovation and push it against other innovative systems where it's going to be a one plus one equals three kind of scenario, and I think that's ultimately what these things are able to do. So, in other words, people can find different solutions that are known to work with your platform and they're able to piece together lots of different parts to solve the problems they're looking to solve, and we're adding value to each other. We're solving problems as a team of technology, not just a single technology platform. So, where can we find more info on Snowflake and the developer platform and the ecosystem?

Jeff Hollan:

My recommendation to folks who are listening, I guess, would be to try it out, I guess if you're keen to try it out, especially since I'm focused on the developer persona, and we know developers love to kick the tires on things as quickly as possible. So, the ones that I would plug right off the bat are quickstart.snowflake.com and also the Snowflake developer YouTube channel. So, on YouTube it's For Snowflake Developers. Those are going to be the best ways to see, “what can I do with this? Is there a way that I can extend or build on top of Snowflake?” If you're more interested in some of the data strategy, data discussions, then just snowflake.com has some great podcasts, blogs, you name it that go specifically deep in that Snowflake world that would recommend folks checking out if there's anything in this conversation that you want to learn more on or double-click in.

David Linthicum:

Yeah, I'm going to watch the YouTube video on the treadmill today. So, where can we find more information about you on the web, Jeff?

Jeff Hollan:

The best place to find me right now is Twitter or LinkedIn. I love to stay engaged and active and have dialog with folks so don't hesitate to reach out @JeffHollan. Feel free to connect. Happy to chat.

David Linthicum:

Yeah, check him out. Those guys have a great data platform, people who are looking to solve issues. So, if you enjoyed this podcast, make sure to like us and rate us and subscribe. You can also check out our past episodes, including those hosted by my good friend, Mike Kavis. Find out more at deloittecloudpodcast.com. If you'd like to contact me directly, you can e-mail me at dlinthicum@deloitte.com. So, until next time, best of luck with your cloud journey. Everybody stay safe. Take care. Bye.

Operator:

This podcast is produced by Deloitte. The views and opinions expressed by podcast speakers and guests are solely their own and do not reflect the opinions of Deloitte. This podcast provides general information only and is not intended to constitute advice or services of any kind. For additional information about Deloitte, go to Deloitte.com/about.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor.

Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

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 © 2023 Deloitte Development LLC. All rights reserved.