



The Deloitte On Cloud Podcast

Mike Kavis, Managing Director, Deloitte Consulting LLP

Title: Leveraging observability to better understand your cloud environment

Description: Observability is the new buzzword in cloud, and many people confuse it with monitoring, but it's so much more. In this episode, Mike Kavis talks with Erez Barak, VP of Product Development at Sumo Logic, about observability—what it is, how it helps organizations deal with cloud complexity, and how it's helped by (and helps) AI. To implement observability successfully, however, a new IT operating model that removes silos and fosters collaboration is absolutely essential.

Duration: 00:23:29

Mike Kavis:

Welcome back to the Architecting the Cloud podcast. I'm your host, Mike Kavis, Chief Cloud Architect at Deloitte, and today on the show I'm joined by Erez Barak, VP of Product Development at Sumo Logic. So, Erez, welcome to the show. Tell us a little bit about your background and what you do at Sumo.

Erez Barak:

Thank you, Mike. Thank you for having me. Thank you for letting me speak on this podcast. As mentioned, I run product development for the observability product at Sumo. Sumo is a continuous intelligence platform bringing in data from multiple sources and allowing for extended, expanded, and highly scalable observability of your applications. In my role, I deal with the development of the product, research of the market, setting the roadmap, and really getting into the nitty-gritty of everything that's observability. Prior to joining Sumo, I spent about a decade in the Azure team at Microsoft learning a lot about the cloud and the needs that relate to data and AI and the cloud, and we see the needs or the cloud needs for multiple customers every day, and at Sumo what we try to do is really help those people migrate and use more and more of their cloud computing.

Mike Kavis:

That's awesome. And that's kind of our topic for today is observability. And before we get too far into it, explain to our audience—some may not be familiar—but what is this new buzzword observability, and why is traditional monitoring no longer enough as we move into the cloud?

Erez Barak:

I think you said it well, Mike, buzzword. There's a bit of a buzz around observability. And, okay, so what does observability mean to us? If I step back for a minute and talk about the environment, so the environment we're operating in is, as I said earlier, the cloud environment. Could be a public cloud, could be a hybrid cloud, could be a private cloud, could be, and most likely, is in most cases a combination of all of those. Within that environment—by the way, not to take away from systems and applications still running in on-prem and having that in the mix, definitely part of the mix. Within that environment, think of the microservices architecture. Multiple, tens, hundreds, sometimes thousands of services running to power your applications and wanting to work together.

And now, think about all the data that's generated across these clouds from the microservices and the complexity it creates. So, when we think about observability, it's thinking about that environment, thinking about the applications we build in that environment, and providing a view into them. We monitor them. We help you troubleshoot them. We help you better understand your applications and improve them over time. We even touch on a security aspect and make sure your application is secure. So, at the heart and soul of what we do, you could definitely think about monitoring. But when people refer to observability, that term definitely brings into mind the environment, context, and scope I just spoke about.

Mike Kavis:

Yeah, so I first started in the cloud like 2007-08ish. We looked at the cloud more as an IS solution, and a lot of clients, when they first went into the cloud, it was move this server to there, so it was still a very traditional—even though we're in the cloud, it's still very traditional NT architecture and traditional monitoring tools did the job, but then as we matured in the cloud, we started taking advantage of that stack. Now, all of a sudden, you have all this elasticity built in, servers come and go, you start pricing containers, microservices. Now, it's incredibly complex. And I think that's where I see this—I think terms like observability, the concepts have always been there, but I think the urgency was kind of born from all this complexity. What do you see?

Erez Barak:

I see similar aspects. I see companies and enterprises being at different stages of the journey, but as you can imagine, it's not clear cut. There are different stages of the journey, and more importantly, at multiple stages of the journey at the same time for their different applications. So, you could see an organization that still has the bulk of their workloads on on-prem, but are taking two, three, maybe a handful of applications and building them as cloud-native applications from day one. You could see an organization that, to your point, has moved most of their infrastructure into the cloud but has not yet switched to more of a microservice architecture and has a roadmap in order to get there.

So, I think you definitely sort of hit it when you talked about the span of stages there. Again, the complexity gets even more interesting as you look at the reality of things and where each one of the applications is, and when we look at observability, we can't just turn away from an application at a certain stage. We've got to be ready to give you that depth of view or vision into the application, or visibility I should say into the application from on-prem, to in the cloud to in the public cloud, to the cases of hybrid cloud, and especially when it comes to applications that move to a distributed microservices architecture.

Mike Kavis:

Yeah, so I was reading a recent article you wrote called, "Observability in 2022: More Open, More Insight, More Collaboration Needed." And there's a couple interesting nuggets I took out of there I wanted to ask you about, and one is you talked about, moving forward, as we build more cloud-native apps or get more mature in the cloud, the business will start expecting, out of IT, kind of insights from how users are interacting with the applications—these are my words, not yours. That's kind of how I summed it up. So, why is that, and how does that change to the roles and the skill sets or the job of the people in IT?

Erez Barak:

First, thanks for reading the article and thanks for calling it out. I appreciate it, and, also, I appreciate the insights you garnered from it. And sort of using your words, I think you mentioned in the context of observability, it starts with the application. It starts with the user. And we know in IT it's not always where you have the deepest view into. We know in IT, we understand the machines we put in place, the servers we put in place, the clusters we put in place, the MKA implementations we have ready to go, the monitoring of CPU, hardware power, et cetera. Definitely the bread and butter of what we do.

But, again, tying back to the evolving world, it's going to be increasingly hard and increasingly challenging to provide value for the users without that tie or—to IT I should say—without that tie to your end users and the applications. If an organization has five applications they are banking on, IT needs to be able to, at the click of a button, navigate from an application view, show me what the user did, show me what screen they looked at or, frankly, on their mobile, what items they touched and where they navigated to, and directly connect that with the things I control with the compute clusters, with the storage clusters, with the network that's bringing in the data, and vice versa.

Show me that network, show me the storage, show me the compute power, and again, with the click of a button, help me tie those to what the user understands, what they support in the application. Now, note this doesn't replace sort of the breadth-first monitoring we do, and the alerting we have in place, and the artificial intelligence we now use to track, and all these services and servers. But that connection helps us really focus on what matters to the business. The business success is tied to the application success. Our ability to tie the work we do from an IT DevOps, DevOps security perspective, back to those applications allows us to connect directly to the business.

In terms of skill set, I think that's where we're trying to really, I'd say, push the limits. I don't expect a new skill set to flood the market. Our expectation is that with given skill sets, we provide the capabilities and applications to support all this complexity. I believe that using an observability application and blending in services, service maps, APM capabilities with traditional monitoring does not require a new set of skills. And the idea there is really extending the power, extending the leverage, and extending the impact IT can have with these observability solutions all the way out to the user.

Mike Kavis:

Yeah, I agree. I don't think it requires new skills. I think it requires a new operating model. The traditional operating model where the app team throws it over the wall to the ops team and throws things over the wall to security team, and there's a very limited view into what the application does—I think if we get to more full stack type teams or platform teams or those types of things, I think the teams get closer to the product, to the business, and could be incented around product and application of a business.

And I think that can drive better behaviors. I think the old model where we specialize and there's a group of people for servers; there's a group of people for network; there's a group of people for XYZ; I don't think that's helpful in a world where we deploy multiple times a month or a day. I think you need full stack teams. So, I don't know if you're seeing a shift in operating models as we start building these cloud-native applications and embrace concepts like observability.

Erez Barak:

I'll take what you're saying even a few steps further and I'll claim that organizational structures will be required to drive collaboration between developers, observability, and security systems—meaning this collaboration is so key to the success of the business that I predict enterprises will put together centers of excellence, will put together teams that constantly work on this collaboration.

And us getting to that organically, enterprises getting to that organically may not be fast enough. May rely on human nature. Human nature does not always drive accelerated change, but here I would say there's going to be such a strong need for this type of collaboration, IT cannot be siloed. Server people cannot be siloed. Infrastructure team cannot be siloed. Application team cannot be siloed. Security team cannot be siloed. So, org structures are going to have to push that forward and really accelerate bringing that collaboration together.

Mike Kavis:

Yeah, I agree, and I see a lot of organizations building teams around reliability, and I see a pattern and an antipattern there. The antipattern is we create a new silo called reliability which creates more checklists, and they're totally detached from products and services and applications. The other one is, for example, a lot of companies are creating this digital experience, which is basically an abstraction across years and years of legacy, but there's this new digital experience for the customer. I see built-in reliability teams focused on that product, application, service—whatever we're going to call it. That is the model that's working because everyone's driving towards the same KPIs, KPRs for that product. But when we just create the next silo called reliability, it's like the antipattern. Another meeting, another checklist, another step, another week before you get things out the door. So, I don't know, what do you see in there?

Erez Barak:

The pattern, the mechanics, that are driving all of this change, I would say, is all around the digital transformation. It's enterprises, companies switching to more and more, to your point, providing digital experiences and transforming what they do to a digital experience for their customers. That is sort of the biggest turn of the crank. Where customers are extremely successful is when they use that, or when they use that mechanic as the primary factor for pushing forward the organization.

And on top of that, naturally, they layer some teams that know more about reliability. They layer some teams that know more about security. They layer some teams that are expert in data collection, but it's not about those secondary verbs, or actions. All these teams are part of the digital transformation culture. When that happens, then you see all these elements synergizing and working towards a common goal. When it becomes all about one of those subareas, your point, if it becomes just about reliability or just about security, then yes, we often see a new silo immediately created.

So, my advice would be for companies take a step back. Think of your digital transformation. Think what that would mean to the end user and the application user, and once as a company we get our head around that, then we start stepping back and saying, "What do we need to make that happen?" If you go through that process, you're typically going to accelerate a transformation. Anything you do that goes against that kind of pattern would create a lot of—a surprising amount of friction very early on as you try to take this journey.

Mike Kavis:

Yeah, I agree with that. I mean, when a team owns reliability or a team owns security, they seem to be the only ones who care about it. But when a product or an application, everyone involved owns it, it gets a lot more attention across the life cycle of software development.

Erez Barak:

Mike, we have many years and many decades of siloed QA teams. They were "responsible" for quality and no one else. Well, guess what? Once those silos started breaking, and quality became everyone's everyday job, quality of services got much better. The excitement about quality started to be shared across the organization. The passion around having quality of services spread around everyone who's touching the service. And I think with observability, reliability, security, et cetera, we've got to think based on that experience and not go through the siloed mode but go directly into sort of layering that in everything we do.

Mike Kavis:

Yeah, I have an old saying, and I'm going to date myself because I use the word pager here, but the best developers are the ones that wear the pager at night, right, because they're getting a call at 3:00 a.m., so they'll probably think a little deeper when they build this stuff. So, the last topic here—and this could be a podcast within itself—but we haven't even touched on AI and ML and the role that plays in observability. So, without stealing your thunder, I'm just going to lay that one out there. How does AI/ML play a role in observability?

Erez Barak:

You're right. It is a different podcast. There's a lot there, but let me try and maybe make a statement and try to support it with a few sub-statements. So, I believe that—or I will claim that—observability must extend to artificial intelligence. Observability must extend to artificial intelligence. Let me explain a little bit what that means to us. If you look at different dimensions of how AI can help observability, and how observability can help AI, a few things quickly emerge.

One is there's going to be—already there are many and there're going to be many, many more AI models deployed in production. So, you want to extend, you want to observe what's happening in production, you got to have the software, you've got to have the services understand what a model is, understand what the set of features are, understand when data drift is happening and react to it in order to observe the overall production system. So, this new area, this new set of entities now coming into production need that coverage.

Second is AI—sorry, observability, if you think about it, terabytes of data come in and petabytes of data coming in as we speak and getting collected. There's no way to keep up without applying artificial intelligence to detecting anomalies, mapping entities, and driving patterns from the data coming in. So, think of this as more of an enabler for this crazy scale and complexity we've talked about, and you can find in Sumo, and in many other services, AI built in to help you sift through logs, to help you, as I said, find anomalies, to help you quickly do root cause and find or analyze what the issue is, and do that in an automatic fashion. I think, or I believe there is a certain pattern we can follow to help this as it develops in the organizations.

When you think about them in open telemetry, it's interesting to see the work being done there and how that is becoming more and more pervasive anywhere. I believe the open-source community will work on bringing observability into AI over time, and the community is where we're collecting the right signals on which direction to take in terms of how we observe AI and in terms of how we leverage AI within the observability solution. So, there's really no way—and this is a very positive statement—but there's no way around extending observability to artificial intelligence, but the other way to think about it, it's really exciting to see what can happen once you apply artificial intelligence to this idea of reliability in your service.

Mike Kavis:

I agree with that, and one of the other things I always say is, as we talked earlier, the systems are getting more complex, they're getting more distributed. Humans don't scale to watch that, so that's where, to me, AI/ML comes in is the systems need to tell us when we need to look. There's just too much for humans to look at.

Erez Barak:

And people listening to this and everyone, let's not get confused. This is not about AI is going to do this instead of the dedicated SREs that we have working on the problem. Not at all. I think you hit the nail on the head with saying this is about scale. This is about coverage. This is about taking those mundane daily repeatable actions, applying AI on top of them, and then driving SREs to a new level, then driving SREs into the world of insights, driving SREs into the world of being able to look at one dashboard that now covers hundreds of microservices and getting real value out of it. That's what this AI leverage is about.

Mike Kavis:

That's a great place to end right there. Appreciate your insights today. That's it for our show today on Architecting the Cloud. Let me ask you, is there—do you have a Twitter handle or are you on LinkedIn, or where can we find you and these articles that you write?

Erez Barak:

I'm mostly active on LinkedIn and will add the E Barak on LinkedIn. I will add the link there, and I'd be happy for people to connect, ask questions, and follow and get more detail and more content as it comes out.

Mike Kavis:

We'll do that. We appreciate your time today. You can learn more about Deloitte and read today's show notes. Just head over to www.deloitcloudpodcast.com. You'll find more podcasts by myself and my colleague Dave Linthicum just by searching for Deloitte On Cloud podcast on iTunes or wherever you get your podcasts. I'm your host, Mike Kavis. You can find me on Twitter @madgreek65, or just e-mail me at mkavis@deloitte.com. Thanks for listening. We'll see you next time on Architecting the Cloud.

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.

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 © 2022 Deloitte Development LLC. All rights reserved.