



## The Deloitte On Cloud Podcast

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

**Title:** Observability isn't a bolt-on; it's an integral part of the cloud ecosystem

**Description:** Observability is a huge help to organizations in solving systems performance issues, but some still see it as a bolt-on. In this episode David Linthicum talks with Dynatrace's Michael Allen and Deloitte's Jay McDonald about why that isn't so. The trio also discusses why companies should view Observability as a holistic process that can be built as code into every facet of their cloud ecosystem—from infrastructure, to apps, to security, to governance, to compliance, and more.

**Duration:** 00:26:44

**David Linthicum:**

Welcome back to the On Cloud podcast. Today on the show I'm joined by Jay McDonald and Michael Allen. Jay is a managing director and co-chair for modern delivery at Deloitte, and Michael is Dynatrace's worldwide VP of alliances and partners. Good morning, guys. How're you doing?

**Michael Allen:**

Hey, good morning, David.

**David Linthicum:**

So, Michael, I'm going to go to you first to kind of get your story in terms of how you came to Dynatrace. And, so, love to—catch us up with how you got into this particular area of the market. Have you always been working with product companies? What's your history?

**Michael Allen:**

Yeah, I started off actually in the networking space, kind of in leadership positions there, and I moved up to Dynatrace 24 years ago. Started off in EMEA European leadership positions, and for the last five years have been running the partner business worldwide, prior to that running the European partner sales organization.

**David Linthicum:**

Wow. So, Jay, You and I both work for Deloitte, but I'd love to hear your story how you came to Deloitte and what you've been focusing on for the last five, ten years.

**Jay McDonald:**

Yeah, well I've been in IT for 24 years and came to Deloitte about four years ago, coming up on that anniversary. And I've been in the cloud since 2010, so kind of an early adopter. Kind of go back in time, I cut my teeth in ~~package~~ [packet] sniffing technology. I was a firewall developer. And it's been a wild ride from doing operations and into architecture. Here at Deloitte, I lead the cloud engineering practice, what we call modern delivery, and that includes DevOps, Agile, Site Reliability Engineering, cloud native development, and of course Observability.

**David Linthicum:**

So, let's kind of set the stage, and Michael, I'm going to go to you first. So, we have the changing dynamic of IT. We're moving into complex distributed architecture is the way I like to look at it. So, it's hybrid and multi-cloud.

We're dealing with the growth of heterogeneity, we're dealing with the growth of multiple services, we may have had 1,000 cloud services under management five years ago, now we have 10,000 cloud services under management, and, by the way, we need to put our existing legacy systems underneath this umbrella because those are aging and those need to be managed and operated in some way, shape, or form as well as edge computing and IoT. So, we're moving to this dynamic where things are getting harder and harder to operate. So, we're trying to weaponize technologies and tools and approaches to make this happen, and so we have AIOps, and we have the rising science of Observability, and we're going to discuss that. So, how do you think we're dealing with the market these days and what do you see as being in the market that's impacting the way in which we're operating these various enterprise systems?

**Michael Allen:**

Yeah, I think you said it perfectly. The apps are being refactored into hundreds of thousands of moving pieces, being hosted inside these containers where those container platforms and the workloads are kind of moving and resizing automatically. And then you've got the velocity of the kind of release cycles of these apps and services have gone from months and weeks to releases in some organizations that are every day, and it's a connected ecosystem of apps that are running across these multi-clouds.

And I think the challenge is user expectations as well. Everyone expects nothing less than always on. And the challenge, I think, with basic Observability, to your point, that a lot of organizations are struggling with is just taking logs, traces, events, and metrics coming from a myriad of disconnected kind of sources, putting them in a database, and then really relying either on data on glass, and humans to be the intelligence layer, or very rudimentary AIOps systems that try to handle this sort of data overload, but they're basically no better than time-based correlation engines, and they really don't work at scale. They lack context, they lack dependencies, and that really causes a problem because, ultimately, where people want to go with Observability is they want to automate off it. They want to self-heal, but you can't self-heal on false conclusions, and that's the real challenge to that, I think.

**David Linthicum:**

Yeah. And go to you, Jay. So, in other words, we've had self-healing, we had the ability to do monitoring and management, and may have kicked off some events. Observability kind of takes insights into the next level, and so we're able to kick off actions based on insights I don't think that we had available five years ago. So, how do you see the market evolving, and what do you think are the big drivers right now that's moving people to Observability tools, AIOps, things like that?

**Jay McDonald:**

Yeah, I think Michael set the foundation brilliantly. I mean, the volume of data, the variety, the velocity of data is really becoming too large for humans to really analyze, and, so, we need a better solution. And, so, I see that the emergence of AI and ML becoming really part and parcel to the next generation of our Observability capabilities. And I think that's what we need to really start to harness because the complexity, as you mentioned earlier, is only increasing. Distributed applications are becoming the norm, and humans are just not designed to—excuse my technical term—but ~~grep~~ [grep] manually all these logs. We need a better way, and so AI and ML really helps us automate that. But, also, what Michael mentioned is really contextualize it. And I think that's a huge opportunity. It's a contextual data, tying metrics, logging, and tracing to business outcomes.

**David Linthicum:**

So, Michael, we're hearing this topic a lot, Observability, as related to operations, which you primarily focus on, I believe, and, also, security and even database administration, things like that. So, why are we focusing on—I think Jay just summarized it very well—we're not looking at tactical data anymore. We're not just looking at raw data logs to make decisions as to what we should do to tweak systems, better performance, autotuning, the ability to self-heal, things like that. But Observability really takes our insights to the next level. As you say, instead of just looking at data in glass, we're getting the ability to get down to a single source of truth for what actually is going on holistically within the systems and how to act upon it. So, what other trends are you seeing in the market, and is this something that actually is taking off within the enterprises? Do they understand what Observability is?

**Michael Allen:**

That's a great question because I often get asked isn't Observability just monitoring? And I think monitoring is something that worked fairly well in a time where they systems were fairly static, you're kind of putting data onto dashboards and a human's looking at the dashes, and then obviously reacting and responding to that. Observability is—especially when you couple it with AI, AIOps and kind of build that in – it can kind of deal with the complexity of the modern cloud and the velocity and the dynamics of the modern cloud to try and not only detect issues but also potentially predict them. If we've got a world where user expectations is everything, has to be always on, we need Observability to go beyond just providing answers but also to do autonomous cloud, to do self-healing. So, I would say in conclusion it's no longer a bolt on an afterthought, “oh, we need to monitor this,” but it needs to be built in if it's going to deal with this hyperdynamic infrastructure. And that's where I see a big trend with my customers going to Observability as code.

**David Linthicum:**

So, Jay, Michael just talked about how his customers are looking at Observability, and I suspect they would. They're looking at Observability tools, AIOps, the ability to have AI-powered insights into these various systems. What do you see holistically as the role of Observability within the enterprises? How are they taking off, generally speaking, not around a particular technology? Do they understand it? Do they understand that they're building into the planning and operations thing, or is this something that we need to teach them?

**Jay McDonald:**

Yeah, I think you're striking at the heart of what we're trying to educate our customers on is that Observability is no longer a bolt-on. It's not something you think about after you architect your products, your applications. You really need to be doing it up front in the architecture, and it should be right in the bullseye of all the other parts of your software delivery lifecycle. So, we think about it in the context of using DevOps and other practices around setting up your pipelines, your building automation, and you should be implementing Observability into that pipeline early into the code.

And I think Michael just mentioned Observability as code. One of the newest kind of capabilities out there is monitoring as code, or Monaco, for short, and I think a lot of people are latching onto this. We're doing it with open telemetry, we're doing it with Dynatrace, we're doing it with some of these great products in the market that enable us to hit that telemetry much earlier. So, this gets into operating model and teams, like who's using it, whether that's site reliability engineers or DevOps engineers. I think it really becomes a tool that everyone can use, including developers, including executives because data can be captured early on and prevent those issues upstream down the line before they happen.

**David Linthicum:**

Absolutely. So, Michael, in my illustrious career, we've always been talking about dealing with operations at different levels of sophistication. We didn't really make much of a move, but what I noticed over the years is that we focused on infrastructure. So, in other words, I'm going to focus on compute, I'm going to focus on storage, and not really where the rubber meets the road in terms of the value that IT compute and storage brings, and that's the applications. The ability to have visibility into these things that actually drive my business is probably as important, if not more important, than really just kind of monitoring the infrastructure around it. So, I see lots of Observability systems that are moving into application-level operations and monitoring and self-healing, which is a good thing to do. Explain what that is and why do enterprises need to consider that as kind of taking their CloudOps to the next level.

**Michael Allen:**

With Dynatrace, what we refer to that as is full stack visibility, the ability to be able to have the context of dependencies both right from the infrastructure and the infrastructure-as-a-service up to the applications, the Java, the .net, there's a code level tier, into the services to the applications. So, that's kind of what we sort of from IO and infra, all the way up to the application and then be able to look at that all the way from kind of horizontally from user across the internet and through the various cloud instances, so that kind of user through data center view. And, so, you've got kind of that east-to-west and north-to-south kind of view because the app's being deployed intrinsically today like operations as code with the infrastructure.

And the infrastructure's highly dynamic, and the applications are under highly dynamic workloads, and that's really where we're to be able to spot problems, really needs that full stack visibility. And if you try to do it with two separate tool sets which are not running with the same context, it makes it very, very difficult for an AI engine to actually come to deterministic analysis that can be automated upon. Same is true, as well, when you think about automating the DevOps pipelines. It's not only being able to do autonomous operations, but it's also being able to automate the DevOps pipelines and speed up innovation, to speed up high-quality apps that move into production.

And the other piece I think which is really important, we've touched on it briefly here, is Observability is so adjacent to the security market. So, we should not only be observing from performance and availability full stack, but also security for vulnerabilities and to protect the application. That's kind of also that built-in versus bolt-on.

**David Linthicum:**

So, Mike, I'm going to back to you with a follow-on question. You just said something very profound that I see emerging in the market, and that's the ability to kind of link these tool sets and link these capabilities together, either within the same product brand name, but the security should depend on operations, and operations will depend on security. The reality is that many instances breaches that are occurring are going to lead to process or saturation, ransomware attacks that take over storage systems. They're going to be indicated from operational indicators. The ability to leverage Observability to see the current state and how we're kind of taking—what's occurring in the current state and our ability to diagnose an issue and then take care of the issue using automation. So, how are we thinking about this? Is this something where we're pushing two products together or two concepts together? Is this something that's basically merging into a single concept of Observability but moving it into holistic Observability?

**Michael Allen:**

Yeah, it's—for me, it's moving into a single Observability source of data that both looks at the data and the metrics and the health as well as the security from with a single agent or a single source of data. And to be able to do that, it's moving beyond kind of the old world of security, which kind of like protect and perimeter is really very, very difficult to do in a modern cloud where you've got hybrid components are running, many places the velocity of development, things are being deployed all the time, so you need something which is deployed with the app to see where the vulnerabilities actually are, and then to help organizations know is that particular component actually vulnerable externally or is it—sort of prioritize the vulnerability assessments as well.

**David Linthicum:**

So, Jay, same question to you but what are you seeing more holistically? In other words, across the enterprises that are looking at AIOps, CloudOps, the ability to deal with complexity using automation and the ability to have Observability to kind of take a lot of the insights to the next level so we can take actions to build a better infrastructure and operate a better infrastructure? Are we seeing the convergence of operational Observability and security Observability? Or is this something that's just basically merging into one concept?

**Jay McDonald:**

I think we are. I see a collision course happening, particularly in the logging realm. When you think about structured logs, those typically come from a myriad of systems that are being emitted, and whether that's user queries or debugging logs or security logs, there's really no sense in having these separate. So, it makes more sense to kind of aggregate these and put more contextual understanding around them. And, so, I do see a collision course or a convergence happening in the security realm. With that being said, I'll caveat that there's still some separation with the security information and event management systems. So, I think we're on our way, it's a journey, but today where we're at is I definitely think we have the ability to take logging metrics and traces and put them into one beautiful database where we can visualize what's going on. And, so, yes, we are on a path of conversion, but there are some exceptions to that where security has so many dimensions and layers to it.

**David Linthicum:**

So, Jay, back to you, Michael brought up the concept of DevOps and the ability to kind of build Observability and operational dynamics into the application tool chain, so we're doing some testing in terms of how best it's going to be to operate the systems. And if you look at DevOps, it's kind of emerged in some pretty cool ways. Not only we do security testing, we actually test for sustainability operations to make sure the code's written with sustainability in mind so it's using less power. Are we moving to where we're going to be more operationally aware and we're building Observability systems, maybe even Observability as code into the core applications, and are the DevOps tool chains and the developers in general moving in this direction? Sorry for the complex question. What do you think?

**Jay McDonald:**

I do think DevSecOps plays a massive role here. We think about the pipelines that we're building. What we're really doing is we're setting up quality gates along the path and we're setting up thresholds. And, so, the only way to know if we're reaching these thresholds is instrument it and have telemetry, so Observability plays a massive—a huge role, and particularly around compliance and policy. When you're codifying these and you're putting them as, let's say templates, and you're instrumenting them into your pipeline, you want to know if you're in compliance way before you get to production earlier in your user acceptance testing, or your development stages.

So, this is when you establish your error budgets and your SLIs and SLOs and you really want to understand it end to end. So, I think security plays a massive role in Observability into the DevSecOps capability. And those pipelines are truly why developers, DevOps engineers, SRE all need to work together comprehensively so they understand what are the thresholds that we can tolerate along that system and what are the parameters that we can work within without breaking or putting us in a vulnerable position once we really get into production.

**David Linthicum:**

So, Michael, I got to speak at a lot of DevOps conferences and always talk about the ability to build things in the DevOps tool chain and how we build, assemble, and deploy code that are going to be—going to provide better operations after deployment. Observability as code we just talked about that, but what kinds of things do you see augmenting in the typical way that we do development in our ability to get to systems that are more observable in nature and our ability to get to systems that are much more reliable?

**Michael Allen:**

Yeah, I think a lot of that ties back to what we were saying earlier about it being built in, and to do that, it requires Observability that (A) deploys with the apps. That means it kind of auto deploys, it auto discovers what it's been deployed with and kind of auto instruments what it's discovered. And all the baselines, because specifically with the cloud, the old way of doing predictions and capacity planning when you've effectively got limited capacity to a certain extent, the baselining has to change. And then it's also kind of auto detection on that baselining with automatic root cause analysis.

And if I relate to one of my largest customers who runs Observability as code, for one reason because of the scale of the environment, but the other reason because they don't want humans touching and following a dynamic—what is a very huge 50,000-machine-plus Kubernetes environment because humans will make errors. And they've achieved something really, really tremendous. In the last two years, they've doubled the number of customers, their transaction volume has gone up by 480 percent, and they've gone from around 60 hours of downtime to zero hours of downtime in the last year. And whilst the platform and the number of customers has more than doubled in size, their Observability team is the same size. So, it's really allowing their business to scale but keep their operational costs under tremendous control.

**David Linthicum:**

So, Jay, you mentioned something that kind of intrigued me, and the thing is as somebody who's been a developer for many years before I started doing what I'm doing now, having to build many versions of the same application to deal with compliance based on the localization, internationalization, of systems in particular countries and jurisdictions and legal issues, things like that. So, in other words, we have these applications that have to maintain compliance. The old way of doing it would be just to have a one-off version that's really kind of customized for a particular country, but you mentioned something policies as code, governance as code, compliance as code, the ability to kind of do this in a much more dynamic way that really kind of hit me that really should make more sense, save a lot of money, and make this stuff less complex. Tell us about that.

**Jay McDonald:**

Yeah, for sure. I mean, you think about the planning stages of your application, I mean, most of our clients are distributed globally. And, so, these applications need to scale in various jurisdictions with different regulations, different data privacy and sovereignty issues. So, this gets highly complex, and the only way to really handle that is create these policies in code and templatize them so that when you are launching and deploying your application into these regions around the world, these governance policies are already set up and they're already set to the thresholds to that jurisdiction. And this has to happen all the way in the planning process, all the way in the ALM tools, your application lifecycle management. And you really need to think it all the way through to your ITSM tools. And we think about this flow end to end.

And when I say ITSM thinking you're taking these systems because all these things need to be built not only into the Observability but into those tools as well so that you have a full and complete picture because let's just face it these are dynamic environments; the regulations are changing dramatically. And, so, in order to keep pace with that, we need to have an understanding of what—how are we performing in each of these environments. And the Observability tool lets us know, hey, we're breaching this compliance or we're breaching this threshold. And, so, this is where I think the two really come together in a fascinating way in making our lives much better and much easier than it was prior doing this manually with data centers and the legacy world that we all come from.

**David Linthicum:**

So, Michael, let's say we hop in a time machine and we're doing this podcast in five years. What do you think we're going to be talking about as related to Observability CloudOps, kind of taking everything to the next level?

**Michael Allen:**

Yeah, indeed. I think one of the things we transform the way that we build our software, we transform from two major releases of our platform a year to now doing a major release to our near-4,000 customers every two weeks. That's a major release of our platform, and of all of our customers, 90—nearly 95 percent of those customers—are on a release of Dynatrace less than 45 days old. And when we transformed the—we really transformed ourselves to be a—the way we deliver software, a DevOps organization. We've actually gone that one step further to NoOps, so basically operating fully autonomously from—and actually not having any operations staff, just having a relatively small SRE team that builds self-healing into our platform, and I think that's—five years from now, I think a lot of organizations are going to be talking about transforming the way that we operate software to be one of the NoOps style. And same on SecOps to move to no security operations and have that fully automated as much as possible as well.

**David Linthicum:**

Yeah, I think the human participation is going to be diminished over time in operations. We're going to automate pretty much everything, and that's a good thing. Ultimately it's the ability to react to something when it occurs and the ability to kind of take an orderly approach to self-healing, correcting security breaches, things like that. So, Jay, where can the listeners find your work on the web?

**Jay McDonald:**

Well, I'm pretty easy to find. If you go and search DevOps for Deloitte, you should see our landing page pop up. There's framework there where we have some opiated frameworks and architectures that we talk about. One of them is called PACE, which is really the process, architecture, culture, engineering mash-up and all the things you need to do around those dimensions in order to be successful. So, yeah, you can just Google it and it should pop right up.

**David Linthicum:**

So, Michael, same question.

**Michael Allen:**

Well, Dynatrace is all about—we're in the business of helping customers run modern clouds and run modern clouds right, simplifying that complexity, faster innovation, more efficiently, more secure with AI built in really to enable that automation, and you can find more about all of that at [Dynatrace.com](https://www.dynatrace.com).

**David Linthicum:**

This is a very important topic. If you've been reading my stuff in InfoWorld and even stuff we've been talking on the podcast, the ability to do battle with complexity that's pushing up is the single most limiting factor that's going to stop people from being successful—digital enablement and cloud enablement all these things that are occurring. We're just building best of breed systems, and that's going to generate complexity, and we have to have some way to manage complexity, and that's going to be the ability to leverage the concept Observability and then the implementation of how you do it with AIOps, and put this in a pattern where it's controlling things from development to deployment and how we're operating the systems. I can't stress this enough. We've got to get this stuff right, so if you're not focusing on this as an enterprise IT professional, this should be something that's on your short list right now.

So, if you enjoyed this podcast, make sure to like us, 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](https://deloittecloudpodcast.com). If you'd like to contact me directly, you can e-mail me at [dlinthicum@deloitte.com](mailto:dlinthicum@deloitte.com). So, until next time with your cloud journey. Everybody stay safe. Cheers. 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](https://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](http://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](http://www.deloitte.com/about) to learn more about our global network of member firms. Copyright © 2022 Deloitte Development LLC. All rights reserved.