



Architecting the Cloud, part of the On Cloud Podcast

Mike Kavis, Managing Director, Deloitte Consulting LLP

Title: Leverage AIOps to operate more effectively in the cloud
Description: Cloud has changed the landscape for IT operations. Operating models that functioned well in the age of monolithic applications often break down when organizations try to apply them to distributed cloud environments. Fortunately, raft of new operating model theories and technologies has come to the rescue. In this episode of the podcast, Mike Kavis and guest, Splunk's Josh Atwell, discuss some of the new ways of developing and operating IT in the cloud—especially AIOps—and shed light on how organizations can choose the technology that works best for them. One key is matching tools to needs. Another is defining success metrics based on service-level objectives that, themselves, have been defined vis-à-vis cloud-based operations.
Duration: 00:23:52

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://www.deloitte.com/about). Welcome to Architecting the Cloud, part of the On Cloud Podcast, where we get real about Cloud Technology what works, what doesn't and why. Now here is your host Mike Kavis.

Mike Kavis:
Hey, everyone, welcome back to Architecting the Cloud Podcast where we get real about cloud technology. We discuss what's new in the cloud, how to use it and why, but most importantly with the people in the field who do the work every day. I'm your host Mike Kavis, chief cloud architect over at Deloitte. Today I'm joined by Josh Atwell. Josh is a senior technology advocate at Splunk focusing on helping IT organizations evolve to support the growing demands on them. He's a coauthor of several popular books, serial podcaster as I am, and has led numerous technology user groups and is an awarded public speaker. Never known for lacking an opinion, which we will find out here shortly, he guest-blogs on various platforms and Tweets at [@Josh_Atwell](https://twitter.com/Josh_Atwell). So welcome to the show. Tell us a little bit about some of that stuff I just talked about, some of the stuff you wrote, some of the podcasting you do.

Josh Atwell:

Sure, thanks, Mike. I'm excited to be here. Yeah, in my career I've taken a lot of different journeys through the virtualization space, datacenter automation, merging out into the cloud. I worked for the customer side, vendor side, and was involved with producing a variety of different podcasts. But yeah, it's an interesting industry that we're in, and it seems that as all the things evolve and change there's always new and interesting things to talk about.

Mike Kavis:

Absolutely, and one of the things where I first met you was at the first NewOps conference., another great buzzword, but tell us what that movement's about.

Josh Atwell:

Sure. So NewOps is a rationalization and recognition that IT operations has changed dramatically, and transformation is required in order for IT professionals and their organizations to be able to support the needs of the customers and of the business and their own needs for that matter. And while I'm a huge proponent and advocate for DevOps as a practice and a principle, the needs for IT professionals extend beyond just supporting developers. They've got a lot of tenured applications. They've got cloud initiatives, Kubernetes and containerized initiative, backup and networks and VOIP and just a lot of different things that have changed for them. And so we put together NewOps Days and focus our conversation around NewOps to focus on IT Ops professionals and how they can be a contributor and tie into things like DevOps, but also how do they address the multitude of new challenges that are facing them.

Mike Kavis:

Yeah and you mentioned a lot of technologies there. You know, back in the day we pretty much had like a mainframe, an AIX box, or some three-tier application, and I think it was a lot simpler to keep tabs on things, to operate and run things right? And where today we spin stuff up and down we've got Kubernetes, containers, microservices. Everything's getting very distributed and complex which is starting to see a rise of a new buzzword, AIOps, which is kind of – and before I ask you what AIOps is, I created this image and posted it out there to get feedback. That's kind of how our conversation started. And it was kind of showing – I've been in this space for a long time, like yourself – the evolution of operations where traditionally or in companies that don't have mature practices, we have what I call reactive operations where you react to alerts. And usually that means something's broken and there's some customer impact.

And then for many years, many companies moved to more proactive operations, so back in my day, a lot of that was being served as a developer through APM tools. So we would set some targets, performance targets, and when the performance of an API, or some unit of work, moved five percent either direction, we would go address it before the client did. Now we get into what I'm calling intelligent Ops, or AIOps, and there's assisted operations, so the machines are doing the learning. You know, you've got artificial intelligence, but they're making recommendations to us to act. And then you get into augmented Ops where the machines are actually acting, and then to the end of the spectrum is autonomous operations: self-healing, self-running, self-monitoring. Then you start thinking about things like self-driving cars and stuff like that.

So throwing all that out there – it's a lot of stuff, but AIOps, another buzzword – we love buzzwords, and the hard things about buzzwords is they mean a lot of things to a lot of people. So before we get into talking about this stuff, how would you define AIOps to someone new to it?

Josh Atwell:

Yeah, and I think you touch on a lot of important stuff there. The way that we have to maintain application availability and servicing need definitely has changed a lot. And one of the things that's interesting about AIOps as a term – so it was first coined by Gartner back in 2016, and the thing that's great about it is gives this visual – like, everybody has a sense of what it should mean. And as you said, though, everybody gets to put their own definition. For me, I think the easiest way to outline it is putting more trust in machines to do what machines are good at and allowing—and augmenting—the capacity of humans to do what they're uniquely capable of doing when it comes to maintaining and operating the environment.

Mike Kavis:

Yeah, that's a pretty good way – I'm going to have to steal that. [Laughter]

Josh Atwell:

Okay, well, it's on the recording, so you definitely have access to it.

Mike Kavis:

Yeah, and we do podcast notes, so I'll have it writing, too, so.

Josh Atwell:

Oh, nice.

Mike Kavis:

We're going to focus a lot on AIOps in this talk, and even some talk about operating models. So we were just talking before we kicked this off, and one of the things, when we get a new technology, we tend to think binary, and it's like you've either got to use all of this or none of it. And the answer is always somewhere in between. And I think where we settled is you should look at it workload by workload and figure out what do I need? Do I need just the proactive operations? Do I need assisted? Do I need augmented? You know, what do I need, rather than go out and buy platforms right out of the gate. So the question to you is we've got all these tools. We've got all these technologies. We've got bots. We've got all this stuff. Before we adopt all these technologies, what needs to change? You know, you mentioned operating models, stuff like that, but what needs to change before we start buying tools?

Josh Atwell:

Yeah, I think the most important thing that has to change is that IT professionals – their leadership in the businesses have to get much more adept at defining what their key service level objectives are for the applications that they're supporting, as they continue to rationalize where an application needs to be served, whether they continue to host in house, on premises, some of those tenured applications that never go away, or whether or not they migrate that functionality to the cloud, either at a software as a service offering, lift and shift if you will and just placing that application in the cloud where it's

managed there instead, or refactoring it because of its business value and being able to have that agility. But acknowledging that the operating model for each one of those decisions is likely to be different when you look at it end to end, and identifying tools and operational procedures that are a common denominator for each of those. So you're always going to want to monitor logging. You're always going to want to look at metrics. You're always going to want to be collecting data and telemetry and understanding what's happening in the environment.

But with some of those models you're going to be adding a tremendous amount of complexity, and the velocity and variety of data that's going to be coming in to support IT being able to make decisions and understanding what's happening in the environment is going to continue to accelerate and become more and more difficult for the IT professionals to consume. So that recognition and that acknowledgment is like a big first step.

Mike Kavis:

Yeah, I agree, and talking about operating models here, a lot of the companies, especially the born-in-the-cloud, or some of the unicorns are really, really good at this, have shifted a lot of responsibility left, meaning a lot of them have teams that both build and run, which – you know, we hear terms like SRE and those types of things. But a lot of organizations, they're still trying to centralize all Ops, right? We've always had the datacenter team do the Ops. You know, we're going to the cloud, but this is still the Ops team. And I find that a challenge because, like, if I'm doing an IoT app and I have devices all over the place and then I'm ingesting all this data, I think it takes a little different skill set to really be productive in operating that type of thing as a traditional end tier. So at Splunk you guys see a lot of customers. Do you see some kind of abandonment of the central body operating everything and more distributing operations out into the BUs or at least out into the products? Or do you still see people clinging onto the old model?

Josh Atwell:

Yeah, it's very much a little column A, a little column B, and I think that goes back to what you were saying at the very beginning. Like, none of this is binary. It's not black and white; there's a lot of grey. And to be fair, a lot of operational discipline that has been applied for IT operations for the last few decades, right, aren't necessarily applied today because the systems can self-heal in ways where it's not as necessary or there's sufficient resources or there's other mechanisms in the process of that application functioning that can offload some of that. And so I think the big challenge that most of the organizations are facing is, okay, we have a new application that a line of business is developing. They've decided to deploy it in the cloud because they need the agility and the flexibility and the scalability that the cloud provides.

Their core IT department isn't adept at managing applications in the cloud. They don't have strong visibility into the application and how the application has changed, all the services that it's consumed. And so initially the line of business is taking a lot of ownership from an operational standpoint. And not to cry foul on developers, but operational discipline is not the discipline that they've invested in, right? They're developers and that's what they do, and so they're having to learn operational discipline alongside their operations peers who are now having to learn more about application engineering and the frameworks of software and the services that are being consumed. And I'm seeing more and more organizations who are looking at, how do we increase that operational sympathy between the various teams? DevOps has been a big part of helping with that. And then how do we utilize tools that can add value, information, and be part of our processes for both sides, from development and operations, regardless of whether we're residing on premises with our application or the application is being deployed in one of the cloud services?

Mike Kavis:

So do you see any clients changing goals and objectives of traditional Ops? So the reason why I say that is I see too often people see these roadmaps that people like me and you put out – this is where stuff's heading – but they only go there from a technology standpoint. So they follow the tech trend, but they don't change any people and process. So if I'm still incented the old way, does it work? So are you seeing that people are successful changing goals and incentives of people in operations and dev and security along the way? Or is it still pretty much COEs just contributing people to the cause?

Josh Atwell:

Yeah, I think given the rate of digital transformations that falter, or fail, is a leading indicator that change is very difficult. And I know that's a cop-out answer. I respect that. [Laughter] But organizations are going to – when it comes to IT operations, they're going to move slowly. Some organizations are going to move more quickly than others, but a lot of aspects of the business are going to move very slowly. And the indicator that I use, and you know, as I've talked with analysts and then I talk with customers, is when you look at how an IT organization and the business funds technology and when they look at their budget and how they apply their budget to supporting the technology that supports the business, and that is usually a key indicator as to whether or not change is happening.

And some of those things that you'll note are the transition from a cross model to where everything is CapEx to OpEx, where they're moving from taking applications that they want to be out of the business of managing and just consuming as a service, being able to reduce the footprint of infrastructure on their on-premises datacenters, simplify that infrastructure, whether it's defining their own pod model or consuming a hyper-convergent infrastructure, and then looking at when they use the cloud, which services they use in the cloud, how they operate and manage those. And so the answer is you're going to get a little bit of both in an organization, but I think most IT departments, they're having some difficulty in making this transition because in large part of the way they've been measured. And unfortunately the measurements around uptime, your time being failure and your mean time to resolution, are still key metrics.

Mike Kavis:

So what should be the new metrics in this new world, or maybe additional metrics that we should look at?

Josh Atwell:

Yeah, KPI mapping is a really tough discipline. I think a large part of the new KPIs being looked at are around service level objectives and recognizing that the latest generation in applications, they don't necessarily fall flat like your traditional monolithic application would, right? They can have deprecated availability without being fully down. And so a lot of that measurement is let's don't go completely down; let's just make sure we stay ahead of any service degradation, because as a business we recognize that a 20 percent degradation in service yields X amount of dollars lost in revenue. And we've tried to apply that in the past with, like, our web commerce and people abandoning carts and all. And the real key there of course is having visibility. And as we talk about things like AIOps, more intelligent operational systems, a big part of that is being able to look at those environments, look at those service level objectives and defining what the acceptable behavior from the environment is, and then responding quickly in order to maintain it. That's where these systems are going to start really kicking in and coming to play.

Mike Kavis:

Yeah, and I think we should be designing for Ops, right? So we have these powerful tools that can help us do augmentation and stuff like that, but if we're still just building code focused on features and not building for operations, we really can't take advantage of these tools. So how have you seen companies or clients work better together so that they're thinking about not just the feature to the customer, but how do I deploy this and run it better from a developer – how are developers contributing to that?

Josh Atwell:

Yeah, I think a big part of that comes in operations teams providing visibility and tooling or assisting with tooling that lets the development team and the line of business that's leading that application understand the impact that the application's having on infrastructure and on infrastructure spend if it's in the cloud for instance, and providing that data and making sure that they can see that. You know, long-term sustainability of an application far exceeds the amount of time it took to write that code. Every piece of code that you put out there, it's going to stay out there until you change that piece of code. And if you write a function or a service that doesn't change very much from a lines of code change perspective, like, that means the bulk of its time is spent under operational concern.

And there was a lot of jokes this week talking about the unfortunate view of the stock market as things have been reacting to the coronavirus, and somebody made the comment about, "Oh, great, thanks for fixing that memory leaks." Well, memory leaks is a great example of an application being made where – and we've all experience it, right, where a system runs out of memory, there's a memory leak, and it's holding onto memory or it's consuming excessive amounts of disk space, filling up an amount point or a drive. Those things still happen, and they still happen on your next generation applications as well. And by providing visibility to the developers and operationally keeping an eye on those things and then implementing systems who are intelligent enough to see those and notify early that they can see the trend on that hard drive is filling up, those are the key areas where from an operations standpoint become exceedingly important.

Mike Kavis:

So when I was walking through the kind of five phases or five types of operations, one of them that we talked about before we jumped on here was augmented operations, and I think that's the one that made your ears pop up more than the others. So give us an example where you see the benefit of an augmented operation? What is it and where have you seen that applied as real beneficial?

Josh Atwell:

Yeah, sure, absolutely. So as an organization DevOps out shared services, either APIs that serve multiple applications – think of maybe an inventory or customer database interaction from an API standpoint. As their services are delivered out, they become key to multiple functions of business. And what's nice is that like most computer systems, software systems, they can become fairly predictable over time, right? They have their own cadence, if you will, and it is something that over time you can identify and get a feel for, like, what's normal behavior for an application, how it's going to respond during seasonal peaks, how it's going to respond when there's a snow day in a region. Like, you can – or the situation that we're dealing with around the world with more people working from home. What's important with augmented operations is having the platform be able to see when an anomaly is presenting itself anomalous behavior where you anticipate specific patterns within the application and in the systems that are supporting the applications.

And then there are typically leading indicators that something is going wrong or starting to present itself as going wrong. Humans don't always do a great job of seeing the leading indicators, right? We're really good at seeing spikes, either up or down. We're good at seeing color variation being presented to us with green, yellow, and red. But being able to look at millions of transactions, millions of logs, like this large volume of data and correlated data between multiple systems we don't really have the capacity to do that in flight, in real time. So we require tools to help us do that, and that's when the augmentation comes in, because as these systems can identify trends or anomalous behavior or something that the system either recognizes and knows who to contact to resolve or sees something that this is looking bad based on something I've seen before; somebody needs to take a look at this, right? You're now empowering the IT operations professionals and the developers to get ahead of potential service degradation or an outage.

Mike Kavis:

Yeah, that's pretty powerful stuff. I remember in my day you'd have your big dashboard up with the health and then you'd have, like, ten tabs open. One's running a collect. Another one's running a trace. And then you're CADing some program you wrote, piping it to – and it's like you just have to be looking at the right thing at the right time to even have a prayer of seeing that something's going south. And as we move towards highly-distributed, complex applications, it's just almost humanly impossible to be proactive without tooling, is kind of my experience. I don't know if you've seen the same thing.

Josh Atwell:

Yeah, absolutely. It is humanly impossible for someone to be able to track all of that, particularly to track it 24/7, which is – you know, all of our applications are always on, right? There's always somebody looking for an Uber ride. There's always somebody making a bank transaction. There's always something doing something. Like, all this stuff is completely on all the time, and these platforms can be on all the time as well.

Mike Kavis:

Yeah. We could go on and on, and this is a topic near and dear to both of our hearts. So thanks for coming out today. Appreciate your time – cool conversation. That's it for Architecting the Cloud. Josh, just repeat your Twitter handle to everyone. I highly recommend they follow. You're doing some good stuff over there, yeah.

Josh Atwell:

Thank you very much. Yeah, so it's @Josh_Atwell, and I'm blogging again at Josh-Atwell.com.

Mike Kavis:

Cool. Thanks again.

Josh Atwell:

Absolutely.

Mike Kavis:

To learn more about Deloitte or read today's show notes, head over to www.DeloitteCloudPodcast.com or you can find more podcasts by me 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 and you can contact me directly: MKavis@Deloitte.com. You can always find me on Twitter, @MadGreek65. And we'll see you next time on Architecting the Cloud. Good stuff, man.

Operator:

Thank you for listening to Architecting the Cloud, part of the On Cloud Podcast with Mike Kavis. Connect with Mike on Twitter, LinkedIn and visit the Deloitte On Cloud blog at www.deloitte.com/us/deloitte-on-cloud-blog. Be sure to rate and review the show on your favorite podcast app.

Visit the On Cloud library

www.deloitte.com/us/cloud-podcast

About Deloitte

The views, thoughts, and opinions expressed by speakers or guests on this podcast belong solely to them and do not necessarily reflect those of the hosts, the moderators, or 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 © 2020 Deloitte Development LLC. All rights reserved.