Architecting the Cloud, part of the On Cloud Podcast

Mike Kavis, Managing Director, Deloitte Consulting LLP

Title: **Leverage human-led AIOps for more effective cloud security**

Description: With a shift to cloud, it’s crucial for organizations to re-evaluate how they approach security—specifically how they evaluate risk. It’s always been critical to take a proactive, predictive approach, but with cloud that criticality is magnified because of the distributed nature of cloud computing. In this episode of the podcast, Mike Kavis and guest, Trend Micro’s Mark Nunnikhoven, discuss evaluating risk in the cloud. Mark’s take is that, while cloud is more secure because of its shared-responsibility model, predicting the likelihood and type of security events is still incredibly difficult. He also argues that, even with the application of AIOps to bolster predictive capabilities, humans still need to interpret model results and make final decisions on which actions to take.

Duration: **00:23:46**

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. 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 the Architecting the Cloud Podcast where we get real about cloud technology, We dicuss what’s new in the cloud, how to use it and why, but more importantly we do it with people in the field who do this work every day. I'm Mike Kavis, your host and Chief Cloud Architect over at Deloitte. Today I am joined by Mark Nunnikhoven, a security educator and advocate. Mark is Vice President of Cloud Research at Trend Micro. He spends his time looking at how security needs to be modernized, and the cloud is a fantastic change to correct a lot of longstanding problems. I totally agree with that. Matter of fact, I'm going to ask a question on that. We got connected because I watched the AllTheTalks Conference and I actually caught midway through yours. So when I re-watched yours I caught the nice analogy you used on security risk. But it was a good talk, really simplified for people who aren't focused on security like me, the architects and the ops people, put it in English for us. And the discussion was, “Risk Decisions in an Imperfect World,” and I thought it'd be a good topic for today. So welcome to the show. Do a quick intro and then we're going to dive right into it.

Mark Nunnikhoven:
Yeah, thanks for having me, Mike. I really appreciate it. Yeah, as you said in your introduction of me, I look at how security can really be better in the cloud and I think having been a security professional for way, way too long now, 25-plus years, there's a lot that we've done that makes a lot of sense if you follow the trail. But coming into it now, especially in a cloud world, it doesn't make any sense. We need to radically update how we approach security and that's on us, the security folks, not everybody else. And so this talk from AllTheTalks online was looking at one of the most fundamental pieces of security and that's how to evaluate risk - because I think it's horribly misunderstood by most people outside of security and a lot of us in security as well.

Mike Kavis:
Well, that's quite refreshing, someone in security admitting they have a problem, right? [Laughter]

Mark Nunnikhoven:
It's the first step, and it's a big problem.

Mike Kavis:
The first step. "I have a problem." No, but seriously you opened that conversation with a talk-that-I could-explain-to-my-mother type of what is risk and how do we assess it and I thought it was excellent, especially for people listening who aren't focused on security. So why don't you do the Reader's Digest of that? But I thought it was excellent. And then we can go into how things are changing in the cloud.

Mark Nunnikhoven:
Sure. So risk is really straightforward when you stop to think about it. And it's really kind of asking, “Is this decision that I'm about to make going to haunt me down the road?” Right? That's all you're worried about, and that really boils down to two fundamental elements for risk. And the first is what's the impact of this event? And what's the likelihood that this is going to happen? So it turns out security is reasonably good at figuring out the impact of an event. So we see this – you know, if you're in the Windows world or have ever lived in the Windows world, we see this on Patch Tuesdays. There's a new patch that comes out and everybody looks at it and says, "This is really bad," or, "This is okay." You know, "Here's the possibilities if somebody took advantage of this issue." So we can figure out the impact reasonably well based on experience.

What we're really, really bad at is figuring out the likelihood that that may occur. So there are other examples, non-IT related. So if you think about like car insurance. Nobody ever wants to think about car insurance, but the amount of data that insurance companies have to make a decision as to what your insurance rate is going to be, whether they'll insure you at all, is pretty significant and that data helps them not only evaluate the impact – so if you got into an accident, they can figure out, you know, ballpark how much they'd be on the hook for – but they've got mountains of data from decades and decades of analysis to tell you the likelihood.

It's terrifying how accurate they can be in the likelihood of you driving around that car in this neighborhood are 28.9942 percent likely to get into an accident that's going to cost us this much amount of money. In security, we don't have any clue really around the likelihood rather than we just look kind of into the wind and hope, and it's reflected in our tooling, and that makes it really, really frustrating. But in a nutshell that's really what we're trying to do, is figure out the impact of an event and the likelihood that it'll occur, and then we react to that as an organization. Are we okay with that? Are we not okay with that? Do we want to do something about it? That kind of thing.

Mike Kavis:
So we have this imperfect science and then we throw cloud – boom. Now you've got cloud and it changes everything. So in your approach to assessing risk and acting on it, what's still the same when we talk about cloud, but what's different?

Mark Nunnikhoven:
I would say what's still the same is we still have a reasonably accurate view of an impact of an event. So when we see a vulnerability, which is basically just a mistake somewhere in the code, or a misconfiguration, we can kind of figure out the impact. So a really easy one that's unfortunately all too common, is if somebody has misconfigured the permissions on something like an Amazon S3 bucket. You can figure out the impact of that event pretty straightforward. You look at the data that's in that bucket and you go, "You know what? It's the lunch menu; the impact is pretty much nothing." It's 200 million voter records with personal information? The impact is significant. There's going to be lawsuits. There's reputational damage. We can figure out the impact pretty well, even in the cloud. It may be a little more difficult to keep up, but the process is still there. And the likelihood now has gotten even harder to figure out, though, because now we're in this highly-dynamic environment. We've enabled development teams to have this great amount of power and tools to innovate. But that means they're doing all sorts of crazy stuff, which is great, but now it's even harder to figure out, "Is that bucket getting misconfigured?" We're not sure.

Mike Kavis:
So – I know in the cloud there's a lot of continuous security, continuous compliance, kind of always running stuff. Is that really new, or were we doing that stuff before cloud, or just not as much?

Mark Nunnikhoven:
We tried to do it before cloud. [Laughter] It's not a new idea, is probably the happy answer there.

Mike Kavis:
Okay.

Mark Nunnikhoven:
The challenge we've had is security is traditionally in firefighting mode all the time. We're always trying to put out fires around vulnerabilities, around attacks, around issues, data breaches, stuff like that. There's always something that's gone wrong that security is responding to. The challenge you get in a traditional environment is that all these systems, there's no common interface, so it was really hard to do any level of automation. You had to go out to the teams and kind of beg and plead and convince them to use the tools that you were providing, if you were at the point where you could actually provide them with something. When we move into the cloud, most companies, most organizations, even if they're multicloud still have the lion's share of their deployment in one environment. They're in AWS or Azure, GCP 80 to 90 percent of their stuff is in one place, which means there's a common set of APIs, which when security teams can take advantage of, is a huge upside, because now we can automatically scan for this stuff and stop pestering people. And it gets to the point where security can take action, as opposed to hoping somebody else does, or negotiating with them so that they take advantage of these tools.

Mike Kavis:
I work with a lot of companies trying to help them transform the way they work as new to the cloud, and there's a lot of shift left – you've probably heard a lot of that. And one of the things we talk about is actually embedding security engineers in Sprint Zero, sometimes permanently on teams. There's the whole security for the business, you know, the common, everyday stuff, but then there's security for this app. Do you see that happening in your line of work, where we're actually embedding security architecture like certified in the cloud of choice and basically living in that AppDev world with those teams?

Mark Nunnikhoven:
Yeah. I think it's an interesting model, and it's something that people are going for and nobody's really sure what the right balance is. The thing is that this is not new ground. So it may be new for security, finally getting out our shell and actually coming out and talking to people instead of just being grumpy in our corner, but there are other skill sets that are similarly sort of – there's a lack of trained professionals. You know, there's not enough of them to go around to put one on every single team. UX is sort of the first thing that always comes to mind. So for usability engineers, we can't normally line up one to one for every team in an organization, so we have them cover multiple teams and they sort of float back and forth, and then all the UX folks will talk to each other. In fact, Peter Merholz, who's a UX expert, he wrote a book a few years back around exactly that problem. How do you scale a limited pool of talent to a much larger community? And I think security needs to follow that same vein, that we need to continually talk to development teams. We need to be continually working with them, definitely on Sprint Zero, but all the way through, because you know, just like software quality, security is cheaper and more effective the earlier you do it. But the challenge is there's not enough security people to go around, so we need to kind of scale out like the UX folks do.

Mike Kavis:
Yeah, and a lot of teams in response to the dangers of you know, all the shifting left, they're building what we're calling cloud platforms. So for example, I'm using AWS or I'm using Google. We're going to run that for you. We're going to put some guardrails around it and we're going to manage the account, and then let you come in and use services. So it's funny because, before all this consulting, I did a startup – this is like 2008 – and all our competitors weren't in the cloud and we were. And what was funny was we'd sit down, and we'd have like 20 people, potential buyers, in the RFP process interviewing us, and they'd get to the end and they'd go, "Oh, you're in the cloud. Well, here's 100 more security questions for you." And I said, "Well, so the other guys didn't have to answer these?" So I actually saw cloud as an advantage from a security standpoint, of course if you have good security hygiene, but there's an opportunity to bake a lot of this stuff into those guardrails, which we didn't have that luxury before. Do you see an emergence of platform teams, or do you see – you know, automation's only as good as your automation, but do you see the ability to make things more secure through cloud tooling, automation, all that kind of stuff?

Mark Nunnikhoven:
Yeah. There's a bunch of stuff in there that I absolutely love. So I love the term guardrails, because it shifts away from the traditional model of control and sort of absolutism, where it's like, "No, you're not allowed to do this," or, "Yes, you can do this very narrow vein of activities." Guardrails implies like, "Here's some really bad stuff. Don't fall off a cliff, but other than that it's up to you to drive the car. Go ahead." I love that concept. I think that's where we need to go. But in general, I think you can be far more secure in the cloud than you could in traditional environments, simply because of that whole shared responsibility model with your cloud service provider.

Now culturally, that's really hard for a lot of people in security to wrap their head around, because there's a bit of ego going like, "No, my teams are great; we can do this better." And even if that was true – you know, pro tip: it's not – but even if it was, why not delegate at least half the responsibility to a world-class team at one of the big three? Like, if I can get somebody else to do my work for me at an adequate level, that's a win because that means I can focus on more important stuff, stuff that's closer to the business, stuff that can't go outside of your org, like truly understanding what drives your business and what moves forward. So when you get it, like when teams really understand the cloud, it's a huge accelerator, not just for development, but I think for security as well.

Mike Kavis:
Yeah, and those big vendors probably have a little more budget to apply to that than most teams do. Yeah, I mean, you guys are a security company so it's probably a little easier to get budget for security stuff, but most companies, you know, it's just like you have to scream and yell and claw to get the money to secure those things. So it's like, "Where's your core competency? Where's your value? Right?" That's cool. So we're going to talk a little bit about AI now, right?

Mark Nunnikhoven:
Sure.

Mike Kavis:
So AI's creeping in here. I've done a number of talks about it, but there's levels of artificial intelligence ranging from kind of assisted operations or security, to augmented, to fully automated, right? I think a lot of security folks are just coming around to the cloud. Now we're throwing AI. How is that being accepted? And what's your take on where AI is going and how that can help and actually create more issues than good as well?

Mark Nunnikhoven:
Yeah, and I think the challenge with AI is the breadth of the term, and it's gone through the whole marketing hype cycle, the popular mainstream media hype cycle. Normally when you say AI to people they're thinking like sci-fi shows, like the Terminator and, like, fully autonomous things, and we're nowhere near that. But I think AI in security is a really interesting sort of blend of really good successes and utter and complete failures. So where we've seen some really great success is in things like behavioral analysis for code execution. So it turns out that how code should run in Linux or Mac OS or Windows is relatively predictable, right? There's a set of APIs that the operating system provides for your programs to call based on users actions, things like that. And this models really, really well in machine learning in order to predict like what's a sort of standard application execution versus what's malicious, right? So if I'm typing in my word processor, there's a whole bunch of stuff going behind the scenes that's "normal," whereas if there's malware running in that document, or in a macro in that document, all of a sudden that lights up like a Christmas tree to a machine learning model. So it's a really good success because that's not something that normal signature matching or pattern matching is ever going to find.

Where we're seeing a lot more failures, where people expected to see successes, is in things like event correlation and investigations. And I think honestly the reason there is that humans aren't very good at that; why would the AIs we write be that much better, right? It's connecting these disparate events. And, when it comes to backtracking an instance, you end up normally collecting a whole bunch of what would be deemed low severity events like two-out-of-ten, three-out-of-ten kind of thing. But in context when you're looking back you're like, "Oh, yeah, okay, that's what happened." Like they put the key in the lock, they opened the door, they walked in, they walked out. That makes perfect sense, but individually they didn't. So AI's been tried and there's some great research going on there, but it's nowhere near where people expected it to be and they're definitely nowhere near where a lot of marketing teams claim it to be. So it's a mixed bag, right? There's some good; there's some bad.

Mike Kavis:
Yeah, and I love the concepts of AI. I liken some of it, though, too, like a dependency. So, like, today say you're going to go on a trip 400 miles away. You've never been there before. You just pull out your navigator and you go. And then all of a sudden the power goes out and you're like, "Oh, crap, I didn't plan," right?

Mark Nunnikhoven:
Yeah. [Laughter]

Mike Kavis:
And I've done that. I'm like 400 miles down the road, and all of a sudden there's no connectivity, and I've got 400 miles to go and I'm like, "I never really studied this." So there's a little of that when you become reliant on things. The other example I use is, my kids are in their twenties, but they grew up with cars with all these gadgets on it, right? You don't even have to look at your blind spot. I think if they get in my car they're going to kill themselves, right, because they're so dependent on technology. So, like, when my daughter was first driving I made her drive my pickup because it had no technology and she had to learn.

Mark Nunnikhoven:
Smart.

Mike Kavis:
Yeah, but she still is totally dependent on all the technology now. But those analogies we – there are some cases where – that's where I like the assisted AIOps. Maybe that's where you start, where it says, "You should probably do this, but let humans interpret that." But then there are some things like when

you get to self-driving cars, you can't have people doing that, or you have IoT devices scattered across the planet. You can't send people out there. So there are some cases where you really have to do it. But there's this dependence thing and there's this ethics thing –

Mark Nunnikhoven:
For sure. [Laughter]

Mike Kavis:
Yeah, I just said a lot there. I'm just doing a brain dump on this, but where's your take on all this? Again, we're early in it, but where's your take on all this stuff?

Mark Nunnikhoven:
Yeah, and II think the easiest place to kind of unravel some of that is around that augmented idea, because I like that, too, because a lot of the time when you build – so if you've ever built a machine learning model, or gone into some other areas of AI, like particle swarm optimization, or fuzzy logic, there's a lot of this stuff where you, the developer don't actually even understand what's going on, right? You train it up and you give it the basic inputs and then it matures those significantly to create a model to make decisions based on your eventual inputs, kind of a thing.

Think of it like the game of telephone, you know, with 20 people around the circle. Use the developer to start the message, and you set up the people around because you thought, you know, every third person was smart enough to actually recreate the original sentence, but you really have no idea what's coming out. So that augmented idea is really smart because instead of making that final decision, at least it comes up and says, "Based on these criteria that I've found, here's what I think, and what I would've decided as an AI. But you as a human are either the sanity check or the final decider." Somewhere in there is good.

So it's really that idea of like black box algorithms versus sort of gray or white box ones, and that's where I think we need to be for quite a while because even the autonomous car one – if you think of it from like a safety issue, we as drivers react based on a history of things. So if we're about – ethics always have horrible, horrible examples so forgive me, and this is similar to how I started my risk decisions talk in AllTheTalks as well. But if you were coming up to an accident – you were about to swerve out of control or someone came into your range and you had the option of getting hit or taking out a pedestrian, or a patio of people sitting there eating, you as a human would react on instinct without a conscious thought.

But when we build a car with an AI, that AI is making a decision based on criteria that we gave it, and that's a really, really scary thing. And obviously that's an extreme example, but that rolls into everything. So even if we go back to the ironically simpler, or at least easier to understand or easier to debate, aspect of, like, letting AI deal with security data, at some point it's making decisions that you may not accept, or you may not understand as a business, and you may not agree with, but you've told it way back in the beginning, "Okay, here's what I think; here's my risk tolerance." And it's not updating and changing based on our experience. It's basing it on its own, which is really, really interesting, but also kind of scary.

Mike Kavis:
Yep. So getting back to risk for a second, so a lot of what we talked about is on the receiving end and looking at risk from incidents. Do you see us putting these risk decisions up – so we're talking about AI. Before you even go try to implement this AI use case, do we apply risk assessments to that and think about things like we just did, the ethics of it, the decisions? Or are we too far behind the eight ball and people are just implementing this stuff and we're having to, deal with it later.

Mark Nunnikhoven:
So we should be, is the answer. We should absolutely be considering risk at every stage. The reality is we're not, and so we see that all the time. So you mentioned IoT. IoT is a really good example of bad risk decisions in general, because a lot of the startups in the IoT space are not even the traditional full stack definition of front end and back end. They're doing hardware, backend services, frontend services for apps. Like, they're building so much stuff that they can't possibly have the expertise to figure this out.

So they're making decisions just like, "Oh, it's responding. Cool, we're up and running." And they're creating very insecure backend infrastructure because their decision was, "We just need to get this out the door. It worked. Great." As opposed to sitting there going, "Well, wait a minute, what this actually supports is devices that are going into people's houses to heat them, or door locks," or whatever the case may be where you really should be considering the impact of your actions and making sure that your systems are not just secure but implement privacy by default, all these kinds of things. And it's one of those challenges where on paper, just like architecture, right – you can write out beautiful architectures on paper. As soon as you put business concerns and people against them, things get messy quick. The challenge with risk decisions is that it's always easier in retrospect, but they're far more effective when you think about them before things happen.

Mike Kavis:
So how do you make those risks a first-class citizen before stuff's built? I mean, that's the million-dollar question, I guess, but is that a culture thing or – I mean, sometimes it's to the detriment, right? Some companies, like, you can't do anything unless you come through us, and that's not what I'm asking. But what I'm asking is, like, “All right, we're going to build something; let's include the risk folks to help us make smart decisions.” I don't see a lot of that. I don't know if you do.

Mark Nunnikhoven:
No, I don't, and I've worked in highly-regulated environments. I've worked with organizations around the world over the last few years in any kind of vertical you can think of, and I rarely, rarely see that. And I don't think that model would even work, because again, there's not that many people and there's so much ramp-up time. But I think – and one of the reasons why I gave this talk at the conference was really to help educate people about how those decisions are made, because I think there's a lot of low-hanging fruit. You can make a lot of headway with some basics. So just even understanding that a risk decision is really based around those two aspects, the impact and the likelihood of something, that helps you make a decision pretty straightforward, or at least go out and find some information.

So if we're trying to build something out in the cloud and we're sitting there going like, "Should we encrypt the back end?" Right? It's a pretty simple decision. If you talk to a security person they're always going to say yes. Is that the answer? No, it's not always yes. There are reasons to and reasons not to. But if you don't have that debate, if you don't at least look at the risk and say, "Okay, what's the risk of us not encrypting this, and what's the risk of us encrypting it because there's risk both ways?" You don't have to be an expert to figure that out. And in fact we're not experts in everything. We all have Google. Just Google it and get some basic data, because even a simple risk decision is better than no risk decision and just assuming the risk without being aware of it at all.

Mike Kavis:
Totally agree, great way to close this out. So that's our episode for today on Architecting the Cloud with Mark Nunnikhoven. Mark, where can we find you on Twitter, and are there any websites, podcasts, or anything that you're on that we should go find to learn more about what you're talking about?

Mark Nunnikhoven:
Yeah, I'm on Twitter and most social networks @MarkNCA, M-A-R-K-N-C-A, and my website is MarkN.ca, like Canada, and that's where I publish all my stuff, even if it's sitting on some other medium as well. So those are the two best places to find me.

Mike Kavis:
Cool, thanks again. I enjoyed the talk, enjoyed your talk at the AllTheTalks Conference. 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. If you would like to contact me directly you can always reach me at MKavis@Deloitte.com or you can find me on Twitter, @MadGreek65. And that's our show for today. Thanks for listening. We'll see you next time on Architecting the Cloud. All right, man, I appreciate it.

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.