NETEVENTS

EMEA IT SPOTLIGHT

FINAL

Conference Debate Session VI - AI Meets Security: Next-Generation Tools Leverage Machine Learning for Detection - and Instant Response

Chair: Fran Howarth

Practice Leader Security

Panellists:

Saqib Chaudhry	Chief Information Security Officer, Cleveland Clinic Abu Dhabi
Ali-Reza Moschtaghi	Group Chief Enterprise Architect, Nando's UK
Roark Pollock	Chief Marketing Officer, Ziften Technologies
Andrzej Kawalec	European Director of Strategy and Technology, Optiv

Fran Howarth

Good morning, I'm Fran Howarth from Bloor Research and it falls to me to talk about AI. Now yesterday at any time that AI was mentioned there was quite a bit of sniggering in the room. As you'll see from this, this instant response, no matter how far AI goes, it's incident response, it's never going to be instant response, we're just never going to get there.

Is it hype or is it reality? Everywhere you look in all these news reports it's AI. I'm a bit cynical about this. I think the fuss about GDPR has died down and now people are like what's the next big thing to get behind? I think everyone is now talking about AI. If you look at this definition here, which is a well-accepted one from 1993, that is not really what is happening at this point and I don't see it happening really at any point in the future.

AI has been talked about way back, it goes back to the Greek philosophers where they said in the future this is what is going to happen. Now in the '50s it started becoming an academic discipline and in the '60s it became sort of the stuff of science fiction. You think *2001: A Space Odyssey* and how the computer taking over everything and the A in HAL means algorithmically. It all sort of fell apart in the '80s, where we had the AI winter. It was the trough of disillusionment, funding ceased, and everyone realised that it really wasn't happening.

Now what changed is in the 1990s it sort of became more realistic and then in about 2012 big data started to become a thing. With really big data sets you can't just get anything out of it with humans, we don't have enough skills or enough people to do this. We have to find a way of sorting through meaning so that we can actually get results from it.

Now within AI we have these subsets and I am postulating that what we're really talking about at this point is machine learning. This, as we go back to HAL, algorithms, it's maths and algorithms. It is using computers to sort through the chaff so we actually find what it really represents, so we can detect the threats that are coming in and we can hopefully automatically respond to them in the way that humans are capable are doing but it's all got to be done by algorithms, which have to be written by humans. It's all about statistics and it's all about predicting through the use of maths.

It is coming into use in cybersecurity, it is seen as a better defence than anything else we have, but defence isn't enough in terms of prevention at the bottom of the slide: prevent, detect, respond. We've spent too much time looking at prevent, trying to stop things coming into the system. But we all know that hackers will get into the systems, they'll try to get a foothold and they'll try to stay there for as long as they can. The onus is on us to detect as soon as we can, so that we can respond in the best way possible, automated where possible because we don't have enough skilled humans. It is far too expensive to do this, so this is where machine learning is hopefully taking over.

Now because I'm an analyst, I'm throwing some statistics up. I'm not going to go through these, they are in the handouts that you will get later. I think this one on the left though is interesting. It comes from Forbes Insights, sorry [Davy] and I don't really believe that 92% of organisations have an enterprise AI strategy, because I don't really think that they are talking about AI as it really is. I think they're talking about yeah, it would be nice to have some machine learning, but I very much doubt that boards are taking this as seriously and are actually thinking about budgeting for it.

More statistics. Now, there are two main use cases that I see in terms of cybersecurity at this point. Behavioural analysis, we have so much information about what's going on, events, the users, the endpoints, what they're doing and this all provides context if we can sift through it. Who is doing what, when, what are they doing it and what damage are they causing, how do we prevent that damage before it happens? But it really does take algorithms and that requires humans to write those algorithms, it's not just automated. We need to have visibility over events if we're going to be able to see what is really going on on the network, if we're really going to be able to win the war against hackers and we need to be able to do it faster.

We are also seeing in terms of incident response, now instant response is a panacea. It really is at this point you get security operation centres where you get loads of people who are up against it all the time, they really are trying to fight a battle and it can't be done by humans alone without some kind of automated help. We really are seeing this come in to some of the technologies available. You can get playbooks where you get which responses you should take, and this is only possible if you learn from what's happened before. You forensically examine, you see how the playbooks are running in practice and over time this will get better, so the response will be much more accurate and automated. It is happening now, and people are using it, perhaps it's not really there yet, it's not that widespread as far as I can tell. My panel, I hope, are going to put me right.

There are certainly limitations, dirty data, bad data, bad results. My colleague, Philip Howard, when he talks about machine learning he talks all about governance, data quality, getting your data right in the first place. Security is just a use case for machine learning. If you don't have all that in the background, it's never going to work. You need to take time; you can't just put these systems in and think that tomorrow everything's going to be okay; it's going to be weeks and months and over time they will get better. But that's costly, it takes resources and complexity.

There are so many systems. Yesterday they were talking about how many security systems are in place in the average organisation. Now how do you tie these all in together and actually enable AI or at this point, machine learning? Complexity is a real issue that is yet to be solved in security.

So we're just at the beginning. It's not a silver bullet, there's a long way to go and it will improve over time. Now I'd like to ask the panel, we have a systems integrator, we have a technology vendor and we have two practitioners. What is actually happening? What do you see in your businesses? Where are we? Perhaps we can start with you, Saqib. A lot has been said about the healthcare industry, including with the IoT. You get 40% of the attacks, you have a greenfield site in Abu Dhabi, so you have a chance to put these systems in right from the beginning. What are you doing?

Saqib Chaudhry

Sure, so a couple of years ago we came to a realisation that our SOC analysts, security operation centres analysts are at a disadvantage because they're facing, or we are facing exponential growth in terms of threats, malware. Then millions and billions of events, potentially malicious events then get generated on a month to month basis. Add to that it's just going to become cost prohibitive to just keep hiring more and more SOC analysts.

What we decided to do is to relook at our overall cyber strategy and in a nutshell, the new cyber strategy that we're working on starts from human-based analytics, taking it to the next level which is predictive analytics, so you use big data analytics plus machine learning to predict what can happen. Taking it then to a level above, prescriptive analytics, where we used supervised AI models to be able to (a) predict

what's going to happen and (b) figure out what is the best course of action to do remediation proactively.

Then the last stage is cognitive. It's more unsupervised AI model, where self-learning is expected to happen and then AI using technologies like robotic process automation, to be able to not only predict, detect, but also in the case of a compromise, run playbooks orchestration using RPA.

Now, taking that strategy, we have four focus areas where we see a lot of potential in terms of AI. The first one we see is threat management, so the idea is to use AI to understand or to figure out applicability of threats and also the scope of the threat based on our footprint on the internet, based on our footprint on [prem] and also our weaknesses and vulnerabilities. So match threats based on what we have.

The second area we're looking into is user and entity behaviour analytics, mainly to deal with insider threats, because I think there is different research done. All of them say that more than 50% of the cybersecurity incidents are related to insider threat. The next use case we're looking into is endpoint protection, so we're looking at a couple of technologies, EDR-related technologies that use artificial intelligence to help us prepare for zero-day type attacks and lateral movement.

Then the last use case is incident response, so use again machine learning, AI, to automate and orchestrate incident response, with the goal of reducing significantly the time it takes a human analyst to respond to an incident.

Fran Howarth

Perhaps we could move onto you, Ali-Reza, you're also a practitioner. Does that resonate with what you're doing?

Ali-Reza Moschtaghi

For people who don't know, it's a hospitality restaurant business. So when it comes to IT, we're probably a bit less mature. When it comes to AI or ML, a lot of the industry is looking at that use cases of looking at things like forecasting customers coming into restaurants, looking at footfall, looking at potential - predicting customer behaviours. When it comes to AI, can I help with that? The whole AI and ML or some of the other new technologies are going to reinforce the need for security significantly. But whether that's AI or ML or it just sits on an Excel sheet, which is what previously when it came to sales forecasting was, the issue is still there.

It just reinforces it even more right now if we just have all that data, all that [unclear] that sucks in the cloud combined with Google weather data and traffic data to see which restaurant will have at what time a rush hour and how many people come to the restaurant and how many chickens do we need. So that's the complication, I think to what Saqib was saying, we're by far not there and I don't see we will get there that fast to look at AI, to detect security breaches.

Roark Pollock

The good news for us, even if you don't have a corporate-wide strategy around AI or machine learning or how to implement it, the good news is for the most part almost every security vendor out there today is starting to use some form of at least machine learning in their approach to threat prevention, threat detection or response, depending on what part of the business you're in. You can take advantage of it just with the tools that you already have in place in a lot of cases or putting some new tools into your infrastructure.

Andrzej Kawalec

There's a lovely opportunity that this stuff lends us and as you say, nearly every single security product or technology in there has been given an AI gloss. The IP is being somehow based in machine learning or some secret algorithms that make that particular product so much more effective, efficient or intelligent than all of the other products on there. I think what's particularly refreshing is you talk about having a clear goal to move from human to cognitive decision-making and an absolute need - the burning platform is that security is a stressed and distressed function at the moment, a function that cannot keep up with the volume and the variety of attacks, or new technologies or regulations.

I think people are hoping and grasping for AI as a way to solve that. Yes, there is embedded AI and ML and analytics in every single different class of technology and some work really well when it comes to threat and UEBA is a great use case. In other areas it's just old-school analytics and pattern matching. But when you think about, as we do, when you think about integrating the entire security function, having multiple different analytics AI engines running on different rule sets on different data sets, not being integrated or aligned, means that again when you default back to a very complex fractured technology landscape, it actually doesn't help.

What you need to do is start to put in place that overarching area of strategy and I think focus those efforts around your security operation centre, thinking about how you take the feeds, how you align and understand the validity of the decisions that are being made. Are all these different semi-autonomous decision-making engines making the same decisions on different data sets? Therein lies the issue. I think AI's not making things easier at the moment. We talk about throughput, but actually it should be about goodput. What is this actually doing to help us? We've created another set of talent challenges around data science around people who understand AI and can interpret those things.

We've created another sort of challenge about how to architect systems for an AI-led world. We've created a whole set of new challenges around transparency, how and where are these decisions being made, can we look at the assets? Are AI-based systems robust and can they respond to adversarial manipulation, because of the data in, as Fran said, the data in and the decision-making process informs ever more autonomous reactions and that's the panacea, isn't it? I would really like to step back and let something make these decisions and autoremediate when you see an issue.

But I think we haven't understood the legal frameworks for these things either, the legal frameworks about when a decision is made and autoremediated and how that plays through. So I think there's a huge promise and clear strategies and the security industry is crying out for something which will solve that distressed position it finds itself in, in terms of the human versus machine foot race that we're currently in in terms of volume and the ability to understand big data sets. But when it comes to an integrated set of challenges, I think we're facing real issues. As an integrator, you look at all the different products, all the different functions, all the different use cases and think actually they're all approaching the problem in a different way.

At the moment, I genuinely believe a lot of these analytics, AI, ML use cases are focused inwardly. Actually, a common data set being looking at the external threat and the risk of different systems and have an entire IoT landscape or user base would be really useful, because that would give us a common set of decision-making criteria against which to put our predictive and preventive measures in place. So I think there's a whole new set of challenges that come with this exciting future. However, it is going to be, I think, one of the most disruptive classes of technologies, certainly in the security industry and it will give us, as we go through the foothills and on that journey through predictive to cognitive, as you described, it will give us hopefully a step change.

But in the meantime, we need to find talented people who understand it, we need to integrate all the different technologies that are using it or not using it. We need to understand how these decisions are being made. We need to understand the human interface, the safety interface in autonomous decision-making and how you start to judge that and the transparency of those decisions. So it opens a whole new class of domain specific expertise that we need and I think people are looking at the effectiveness or efficiency of point products, rather than the system-wide implications of relying on large scale autonomous AI systems.

Roark Pollock

There's no question today, a lot of the machine learning and AI that's getting applied into security tools, it's all about helping us get beyond human scale and address some of the issues we have from a data perspective, the scale of the data and the efficiency of humans to respond to that and to analyse it. One of the things that we primarily use it for is to - from a threat detection or a threat prevention standpoint is to get beyond what I call the old-school, the human protection gap. If you think about endpoint security back in the day and still today, a lot of the approach was once something is known you can then create a signature, create heuristics to identify and block it, or at least alert on it.

Well that model, it took two weeks on average to essentially update malware registries, get the signatures pushed out, have the disseminate down to all of the endpoints that you're protecting. So in that time you've got a two-week on average, sometimes much longer, protection gap where we know about this threat but all of your endpoints, all of your devices have no protection in place yet until all that information disseminates down. With machine learning in place, you can shorten that gap to almost instantaneous.

Then you can extend it beyond just the known, because I can use machine learning to then get beyond blocking just the known issues, to going beyond that and applying that and being predictive in a manner so I can now block or at least alert on issues that - based on what the data you've give it to learn, you can apply it and say okay, I'm confident this is unknown but it's bad and block that on an ongoing basis. So it's not just about the scale, but it's also about getting predictive about what you've taught it. Now on the flipside of that, there is a garbage-in/garbage-out problem, so it's all about the data you use to train the system.

Andrzej Kawalec

You're absolutely right, we shouldn't get beyond our skis in the sense that the use of these technologies have completely fundamentally changed how effective our security response is today. You wouldn't be able to do what you do on the endpoint without this stuff. Think about the SIEM engines and function, it is that pattern matching analytics, moving away from manual use cases and rule sets to actually defining them. How specific analysts in a security operation centre can get decision support, can get suggested remediation paths, playbooks tailored to their specific area of expertise. All of those things are moving us a long way along being more efficient and more effective, but we still haven't got to the automated and somewhat more interesting predictive element of things, which I think is a little way off.

Saqib Chaudhry

So completely agree, I think the goal is to have an autonomous immune system, just like we have in our bodies. It's very encouraging, I was reading about a start-up called Mayhem, if I'm not mistaken. They won the DARPA Cyber Grand Challenge, so their product uses machine learning to identify weaknesses and patches them, so that's very encouraging. I know Google uses now a utility which is based on machine learning to identify weaknesses in android systems. AWS has, I think, a system called Macie, if I'm not mistaken, for S3. So all of that is definitely very encouraging for us.

Now from a challenges perspective, I agree with the challenges you brought up. There are a couple more challenges, at least from a healthcare perspective, (a) the incident response automation. For us we have to be very careful because we don't want to bring down or disrupt a system which is directly related to patient care, so we have to be careful.

A cupule of other challenges are we're trying to build an in-house capability, hiring resources. So last year was the first time that we started looking for data analysts, regardless of cybersecurity experience, data analysts and someone who can write our AI machine learning type of algorithm and code. It's very difficult to find experienced resources, so finally we decided to just hire fresh graduates and train them. Also, we're now working very closely with our business intelligence team, because they have data scientists and data analysts to utilise their resources in building some use cases for us.

Another interesting challenge is quite a few of the vendors we're dealing with, they have machine learning, however, their machine learning gets trained within their own cloud

environment. So it's not getting context of what we have in our environment, it's not getting our environment space aligned. So they would run a particular malware in their cloud environment, train their AI and then push it to our system as just another signature. So that's not creating a lot of value, so I think something that we should look into.

Andrzej Kawalec

That was absolutely exactly that point about multiple different systems, cloud-based analytics all using their own AI algorithms give you a more complex and slightly less transparent picture of what's happening than just call logs or incident lists. AI's a little bit like a new puppy at the moment, it's cute and fluffy and everybody wants to touch it, but nobody really wants to take it home because it'll make a mess. It sort of jumps around and barks at you the whole time and you have to teach it what's right and what's not and how to behave. I think it's a little bit like that at the moment.

Saqib Chaudhry

Yes, well considering we see about 111 billion lines of code every year, it soon can be the world needs something like machine learning and AI to go through each and every line of code and identify potential weaknesses. So it's a necessary puppy.

Andrzej Kawalec

Puppies grow up to be guard dogs, don't they, which is always handy.

Roark Pollock

One of the reasons that machine learning is coming to the fore at this point in time though is just this year availability of data that we can use to train the system. There's so much crowdsourced - I'll call it crowdsourced - security data out there with VirusTotal and other reputation services where we have massive amounts of data that have been collected over time, that it's easier at this point than it ever has been to train systems in a way that is much more valuable. Once you put something into your environment, it immediately starts to learn from your environment as well as from what it's seen in the past. So it becomes much more efficient in your environment over time.

Andrzej Kawalec

Your point's right, if not now then when? We've got enough insight into the security function, we've got enough big data and complex things, and we've got some of the tools now that can start to help.

Fran Howarth

We do have a question from Rik.

Rik Turner

I just wondered, this is really mainly for both Andrzej and Roark, have you seen any signs as yet of the threat actors using machine learning or artificial intelligence? I

imagine it's only a question of time because obviously the resource is going to be there, at least from the state sponsored and the guys who have deeper pockets in the threat actor community. Have you seen any signs already in the last couple of years of any AI, ML coming in on the bad side?

Roark Pollock

Yes, absolutely. It's just a tool, it's available to both sides, adversaries and the white hat part of the audience. We just talked about it, even companies are using it today for doing their own pen testing, looking for vulnerabilities in their own systems. It's no different than if you're a black hat or an adversary using those tools for the same purpose. It's going to be used on both used, it's already being used on both sides and I think in the long run, perhaps the best we can hope for is a stalemate.

Saqib Chaudhry

I recently read an article where one of the bad guys baselined the behaviour of an employee who they wanted to impersonate. So they used machine learning to baseline that behaviour. I think if anything, the bad guys will probably make bigger advances than us, because we have a lot of restrictions. We have legal issues to deal with, we have ethical issues to deal with, we have patient safety issues, versus somebody whose only job is to hack. For them it's much easier.

Roark Pollock

We talked about this yesterday in some groups around cybersecurity. One of the reasons the security industry keeps changing and keeps evolving so much is there's no right answer, there's no end game to the industry. We have an intelligent and active and motivated adversary that keeps changing the game. It's like law enforcement, it's never going to go away, it's never going to be a problem that is ultimately completely solved.

Saqib Chaudhry

I can foresee our AI, cybersecurity AI, fighting your cybersecurity AI and humans are just standing back and just waiting to see what happens.

Roark Pollock

Is that a movie?

Saqib Chaudhry

It can be turned into a movie.

Andrzej Kawalec

Just to add to that, we see a lot of AI analytics used to profile and deliver at scale targeted phishing campaigns, for example. So think about adaptive malware, because we've already given the blueprints to our systems and our responses and our regulatory frameworks. So design a relatively simple adaptive piece of malware that knows what

you're going to do and how you're going to do it is not difficult and we've already seen those things. So at each step in that [unclear] chain, whether it's upfront research, or highly targeted at-scale industrial phishing campaigns, or ransomware as a service whereby the ransom that's extorted adapts to the share price of your organisation. These are not sophisticated use cases, but they're just being done at scale.

Roark Pollock

It sounds like marketing, you're just trying to figure out what humans will respond to.

Saqib Chaudhry

There's already malware out there that can detect a sandbox environment. It makes your honeypot inefficient.

Fran Howarth

We do have another question.

Unidentified Male

Joe was talking yesterday around how we're spending 80% of our time on threat detection and 20% on reduction and that really we should be flipping that on its head in terms of where we're investing our time and where we're looking at within the organisation, so that we can reduce threats from coming into the network in the first place or the organisation, the application, the database, et cetera. With the advent of zero trust and software defined perimeters and the ability to have ephemeral networks which don't persist for you to be able to find that, how do you see - and particularly one of the comments around more than 50% of the tax being insider threats, which are very hard to detect against when you trust in your internal network. Do you believe that that is the kind of switch that we need to see? Or do you believe AI and threat detection has a larger percentage to play than 20% of what Joe was saying yesterday?

Andrzej Kawalec

I think you're absolutely right and there are two switches that we need to flick. One is absolutely about a movement away from a fixation on the external threat, to actually being a greater understanding of your own environment. The second flip is with that realisation that your own environment should have zero trust built into it. Once you do that, then you up the ante about the level of inspection, the level of understanding, the level of behavioural analytics within your own environment.

The inside out approach is really the first step to take and if those two switches are flipped, you take a completely different view of security when you're not constantly responding and reacting to external stimuli because you own and control your destiny in a certain sense. The idea that there's a perimeter, or that there are bits of your network that are particularly safe is a fallacy that we need to move beyond.

Saqib Chaudhry

As a security practitioner, I would love to see the attacks are first reduced, but that's not the reality. So security should be a business enabler and business is moving towards connected devices, more information sent on mobile devices. So I don't think that proportion, 80% or whatever threat, we won't - I don't think we'll be able to reduce that. But I think zero trust and a combination of AI, we just have to come up with more innovative ways to protect the device, especially disruptive devices in the age of digital transformation. We just have to think outside the box.

Unidentified Male

I would love to give you a demonstration of how our technology can reduce that attack service for you.

Roark Pollock

Good sales pitch.

Fran Howarth

Okay, well, with that I'd like to thank the panel and hand over again. Thank you.

[end]