NETEVENTS

EMEA PRESS SPOTLIGHT

DRAFT

Round Table II: The Security Professional's Best Friend: Artificial Intelligence

Introduced & chaired by

Rik Turner, Principal Analyst, Ovum

Panellists:

Roark Pollock	SVP of Marketing, Ziften Technologies
Simon Crumplin	Founder & CEO, Secrutiny
Jan Guldentops	Director, BA Test Labs

Rik Turner

This is a security-related session, which explains why they involved me, because although I'm in the infrastructure solutions part of Ovum, technically I spend most of my time writing about cyber-security, IT security, call it what you will. So, we're thinking a little bit about AI and its potential application in the security space, not in marketing, not in CRM, not in any other areas. We're particularly looking at AI in the security space.

It clearly is something that, if you speak to virtually any security vendor for the last three years, they will fairly quickly in the conversation be dropping the AI words somewhere in there. They have to, their marketing teams have told them that they need it in there. It clearly is a buzzword, but we're here to get to the truth, and find out whether there really is any actual meat on the bones, whether there really is any true application of AI in security and, if so, what is it.

So, I thought I'd talk a little bit about where we are now, what's changed in recent years, what might be driving this move to AI, or the adoption of AI in the world of cyber-security, and eventually getting a little bit towards the end of this conversation, a little bit about predictive capabilities, if there are any. So, we'll talk a little bit about that.

Anybody who's worth their weight, their salt, have to be able to tell a story in one slide. This is my one-slide history of cyber security over the last three decades. Effectively, anybody you spoke to in the 1990s, be they practitioner or be they vendor, were talking about prevention. They were able to talk about prevention seriously, preventing the bad guys getting in, preventing malware from penetrating their networks. Their infrastructure could be safe. They could prevent all of those bad things happening, no hacks and all of that good stuff.

History has taught us otherwise, but at least they were able to talk with a straight face about being able to prevent. Over the last two decades, anyway, we have moved increasingly towards a new stance, and nowadays you will find the vast majority, both of the vendors and of the practitioners, the people who are actually having to do the work in the trenches, will say look, the best of what we can do at the moment is to detect once a breach has happened - you've heard all these stories about breaches are inevitable, I'm sure.

So, detecting once someone's in, move to mitigate as quickly as possible, potentially do some damage limitation, do some, whatever, quarantining so they can't run amok within your infrastructure, and then subsequently to remediate, clean them up, get them out, and start again until the next breach.

So, that - you could argue, actually, that that is really a defeat for the cyber-security industry, and I'll pose that to my panellists in a minute. But it is, in a sense. It reminds me a little bit of the people defending the city of Constantinople when it was still capital of the Byzantine Empire, and gradually the folks - gradually, the siege made it through the first outer walls, and drive them into the inner walls, until eventually they breached the whole thing. Notice that we use the term breach. We've adopted it from the world of siege warfare here.

So, one of the things that's change - well, I mean, look, this has being kind, really. 30% to 40% of new malwares are picked up by antivirus signatures, no more than that. To be honest, I think it's a lot less than that. I mean, we saw Symantec, back in what was it, 2014, talking about 45% at that time on the front page of the *Wall Street Journal*. I now think it's between 20 and 30, not much more, across the industry.

The reason for that? Well, obviously we're no longer talking exclusively about the famous script kiddies who were sitting in their back bedrooms as spotty teenagers trying to hack the Pentagon for fun and kudos. They're still very much around, but clearly now we're talking about the criminal gangs in multiple parts of the world, some of them very sophisticated. Equally, the hacktivist community that pops up here and there seems to have periods where they're far more active, and then they go into a lull for a while, and state-sponsored groups, of course, the most sophisticated of them all. Because they've got more money behind them, they can get the best talent and they can spend a long time getting it absolutely right, so they can really target.

So, there’s a lot more malware actors, or threat actors, or whatever you want to call them. There’s also, of course, the active market on the Dark Web for all kinds of exploit kits and other fun ways of getting in so that you can hopefully buy something for pennies on the dollar, do a slight tweak so that nobody’s ever seen it before, launch it at your target enterprise and so forth, and do it very simply and very, very cheaply. So, that’s also another factor, the sheer availability of tools that are on the cheap that weren’t there 30 years ago.

Finally, the cloud that makes it so much easier to go out, rent a few processors from Amazon, and test-drive your new exploit before you’ve even launched it, and make sure it works.

So, there are a variety of factors that just made it so much more difficult for the predictive, the preventative stance, to still prevail in the cyber-security world. This is just a graph that I lifted from a report by the [unclear] outfit with [Kenner]. All that’s doing is saying, look, these are the number of published CVEs since 1999 through 2017. It’s not a linear increase, but you see it went from - I think it was 890, if I’m not mistaken, in 1999. It’s now up to over 15,000 last year. Oh, sorry, it’s common vulnerabilities - yeah, it’s effectively - it’s a public list of vulnerabilities, so that it’s the - yeah. Thank you, yeah. So, essentially, it’s just saying, look, the number of vulnerabilities that are now being published has gone up from 890 to over 15,000 last year, and that only tends to increase.

Equally, the average days from published to exploit, in other words, this - what I’m trying to show here without going into gory details and reading all the numbers is the speed at which an exploit can come out and come at you has grown. In other words, there’s just less time involved to get there, so there’s - so, it’s not only that the volume is growing. It’s also that the velocity of the threat landscape is on the increase.

I mean, people are not - it’s like this. I didn’t - there’s another slide I wanted to put up here, but I didn’t for sheer lack of time. I didn’t want to be too long-winded about this. But what’s interesting is, right, there may have been 15,000 vulnerabilities published last year - [unclear] that runs the CVE list - but it’s, like, 2% of them that are ever actually exploited. So, a lot of the others, they’re never, ever exploited. So, you not only have the problem that you’ve got this huge volume coming, it’s the classic - I’m sorry.

It’s the - you have to keep - people in security always talk about the needle in the haystack, and it’s a horrible cliché, but it is true. In this particular case anyway, the vulnerability space, there’s this vast number of vulnerabilities, ever-greater numbers of vulnerabilities, being published, as that previous graph showed but, by the same token, how do you know which ones are actually going to be exploited and what is the likelihood of that one being exploited, and why do you waste your time worrying about all the others when there are only so many of them - well, 2%. Sub-2%, actually, 1.9%.

So, I mean - and also the speed at which the ones that are going to be exploited are exploited is ever-greater, so you’ve got less time to define which ones you need to focus on. That scenario that I’ve described there, to the best of my ability, provoked a certain amount of responses from the tech industry. Clearly, the tech industry doesn’t stand still. It thought, okay, well since we are moving towards a detection and mitigation stance, vis a vis the previous preventative stance, let’s - there’s a few of them here.

Sandboxing was all the rage a few years ago. It made FireEye a very profitable IPO. Made them for a short while there a major name in the industry. We can argue whether they still are, but certainly sandboxing. By sandboxing, I don’t mean the sandboxing that you could do 30 years ago on your laptop, or on your desktop in those days, but the actual network-based sandbox. So, a great big box, they’d roll it in. they sold lots of them to Wall Street. You’d roll a big box in, put it on your network. Anything that looked vaguely dodgy that you could not actually guarantee was malware, you could put it in there, carry out a controlled explosion, and check whether or not it actually was malicious.

It was very fashionable for a while, they sold a hell of a lot. Then, the malware guys started writing malware that knew it was in a sandbox, and played dead, effectively, or played good, however you want to put it, so that it just was released out into the wilds, because they could perceive no malicious activity from that malware. But in any case, sandboxing is still very much part of the arsenal of the defenders these days, but it’s no longer the be-all-and-end-all.

These are a couple of Gartnerisms here. You know that Gartner gets to christen every new sector in the industry, and so we have the wonderfully known-as UEBA, which sounds like a football association but is in fact user and behavioural analysis. This is quite straightforward. You basically deploy some technology on your network, it watches what you all do, learns what looks like normal for you - well, Rik logs on every day, always does this, that, and the other, he never goes to the payroll database, he never tries to look at what’s going on in finance, he just gets on and does his daily job. Oh, tomorrow he logs on after six months in the company, and suddenly he wants to download the entire payroll database. This looks a bit strange.

So, effectively, you are detecting anomalous behaviour on the network. Often you do it up in the cloud, the actual learning bit. Well, where you do it is irrelevant. The fact of the matter is, that’s what you’re trying to achieve here, is to detect anomalies in behaviour, both of users - in fact, this is the reason - it used to be UBA, but then Gartner went, oh, hang on. There’s a thing here that’s not a user, but it’s also acting a bit strange. It’s an entity. So, it might be your email server suddenly starts to download the entire - tries to download the payroll database. Right, it’s an entity that’s being a bit dodgy here. So, they expanded it to the less - shall we say, less - it doesn’t roll off the tongue so well, does it, UEBA? But in any case, there it is.

Then we have EDR, another Gartnerism. End-point detection and response. Straightforward enough. Put something on the end-point - usually, at least, anyway, most of them, I think, put an agent on the end-point, because of course, end-points are no longer all the time on the corporate network, they’re not on the LAN. I’m virtually never on my corporate LAN as I work a long way from our headquarters. I don’t even open my VPN, but don’t tell my IT department I said that, and I’m away doing all kinds of things with my laptop which nobody ever knows about, so therefore if they had some

EDR they could put it on my laptop and check and see whether or not I'm being infected by whatever I'm doing when I'm on the public internet.

So, EDR. It's interesting with EDR. We'll come back to it later, but it's interesting that EDR initially came along - because you've got all these traditional end-point guys, who used to - the Symantec's of this world, the big beasts who we all know and love that became billion-dollar companies, half a dozen of them probably still today - a lot of those guys, they came along, basically, with a signatures-based model - I won't go into all of that, but you know, the way that works, I think you all know what signatures are, or were - and the EDR folks all came along and said, look, you can deploy us alongside them, and we complement. Then - and we'll make it so much better. We enhance those traditional signatures-based approaches. Then, about two years later, they all said nah, don't worry about that. Get rid of that. We'll replace that with us. In other words, started going for the jugular of the old signatures-based guys, and it upset them dramatically.

But in any case, the fact of the matter is that EDR was another response to that. This is not a tech response, strictly speaking, this next one. Threat intelligence. Threat intelligence is not a technology response. It's an information response. It's just, effectively, I do not know who is going to be - and all of my - let's say I run a bank. My bank branch, I don't know who's going to come in and try, with stockings over their head and a shotgun, and try and rob my bank, so could you please send me photographs of every bank robber that you know on the planet? Then, if you - that's really not even threat intelligence. That's just threat data.

Then, you start saying, hang on, how could you - please, of all those photographs you've sent me of bank robbers, could you also mark with an asterisk the ones that are actually not in jail at the moment? That's kind of narrowed it down a bit. Then, even further, my branch is in - I don't know, Croydon. That's a place in the UK, for anyone that doesn't know. So, could you please also then mark with a double asterisk the ones that are out of jail, and they happen to usually operate in the Croydon area? That starts to become threat intelligence vis a vis pure threat data. But it's a necessary response to the fact that you don't know what is coming at you. You just say, give me more information so that I can at least have a better chance.

Then, of course, along came AI, machine learning. I'm sorry, I mean, everybody talks about - artificial intelligence, we're really only talking about machine learning. I'm - well, Jan would say also that we're talking about some kind of - a little bit more than that, and he can talk about that in a minute, but in essence, anyway - but we started seeing, about three years ago, I suppose, probably - maybe five, I can't remember - but a few years back, we started seeing more folks in the security space talking about the need to apply artificial intelligence in order to do some machine learning.

Basically, what they're doing is, they're doing pattern recognition to be able to see what is going on in the - in the - all of the traffic that is traversing the corporate network to try to see things that might look - that look a bit suspicious, that look dodgy, that look like they might need to be addressed, that they are worthy of further attention. I suppose it's also not surprising that incident response used to be just a practice, just an activity, one of your procedures. Increasingly over the last five years, incident response has also become a tech response, or a technology platform, to enable faster incident response.

Because you're getting more threats coming at you, because you're less able to see beforehand, a priori, to what they're doing, the need to respond more quickly to more threats has inevitably led to a degree of automation of incident response. So, that's also another tech response that we're seeing.

I wanted to say - and this is - I'm throwing this out here, but once you've deployed artificial intelligence to help with detection and mitigation, there are already some folks out there who are saying, well, okay, what about if we start using it for prediction? By prediction, I don't just mean crystal ball gazing, but actually hopefully based on artificial intelligence, the - shall we say this? The promise of artificial intelligence down the road is that it might actually get us to some data science where we can start making some realistic predictions about what is the most likely attack on you, what is the most likely vulnerability to be used against you, and those kinds of things.

Now, I'm not suggesting for a moment that that's where we are today, but that's what people are talking about, and what they are suggesting is possible.

For that, I will now throw it over to my learned colleagues on the panel. Jan, let's start with you. Go on. Yeah, Jan. Let's see a little bit about - Jan, tell us a little bit about your contact with, experience of, AI, and - do you want to start - he suggested at breakfast, actually, that we ask everybody how much they even know what it is.

Jan Guldentops

Who knows what natural language processing is? NLP, right? Okay. Who knows what machine learning is? Okay. Who has heard of deep learning, and really understands what it is? That's the next question. Everybody heard of deep learning?

Rik Turner

Very honest, there. Thank you.

Jan Guldentops

This is the number one problem. Artificial intelligence is the next bullshit term we're using, right? It's IoT, it's cloud, it's something that broad that we understand it, but we don't really know what it's about, right? We've been doing this, especially, for instance, machine learning, we've been doing this in the security industry for 15 years. Your anti-spam is based on machine learning for maybe more than 15 years. So, it's limited technology, right?

The second thing we have to remember is, it's a tool. It is not magic. That's a big thing. If a guy like Elon Musk starts saying that he's afraid that Skynet is going to come along within five years, and the terminators are going to rise with artificial intelligence, we're 20 years, 30 years away from real artificial intelligence, like you would define it. Do you agree, or not?

Roark Pollock

It depends on what you mean by artificial intelligence. I think, to your earlier point - I mean, we use artificial intelligence as an umbrella term, but nine times out of 10 we're talking about basic machine learning, and that's what most of the security vendors are doing today. They're providing some sort of natural language processing or deeper machine learning in their products today. I think the important question is, is to ask ourselves, why are we even talking about this from a cyber-security standpoint? Cyber-security's been around since, going back to your one slide, the 1990s, so we're 20-plus years into this industry. Why has it taken so long? Why are we talking about artificial intelligence and machine learning at this point?

I think really, it's a combination of things. (1) The technologies have come a long way, so that they actually work well. We have access to the data that we need to make machine learning effective from a cyber-security standpoint. The machines that we run artificial intelligence on, or machine learning, if - for us, if you're talking about end-point and servers, I can now run artificial intelligence models or machine learning models on those end-points without bringing that device to its knees, right? I can run machine learning as a security tool on your end-point, your Apple, your servers, your cloud virtual machines, and it only takes up less than 1% of the device. It doesn't kill the device from a processing standpoint.

But I think the most important perspective, and the reason we're seeing more and more, besides the marketing reason, obviously - the reason we're seeing machine learning appearing more and more from a cyber-security standpoint is, our traditional models are failing. I think Rik mentioned it before. Yeah, our traditional models are failing and we all know it. Everybody's talking about - everybody's kind of got to the point where they know, or they assume, that they're going to be breached at some point. It's inevitable. So, now, everybody's talking about how do I focus on detection and response capabilities?

Well, if you go back to, now, thinking about prevention, the obvious answer to a lot of those questions now is, how do I apply machine learning to solving some of those problems? I think there are two very specific ways that traditional security is failing. One is what I call - there's a protection gap. Signature models in AV and anti-spam and all these have been going on for years, and they work. There's a reason they've been effective for a very long time. They work at what they're good at.

A piece of malware comes out, we find it, the big anti-virus companies write a signature to identify every time it gets on your device, and so any time in the future that known piece of malware hits your device, now you can block it. Well, that model takes about two weeks on average, if not more. If you don't connect, and you don't update your signature model, that protection gap can be two weeks or more. So, any time that attack now is out, it's a known piece of malware, I can protect you from it. But it takes a while, and if it's an unknown piece of malware, it's not going to protect you at all until such time as somebody finds it, and then you write a signature and you protect against it.

So, there's a protection gap, or a window of time, where the traditional model doesn't work.

Also, most of our traditional security models are built around a file. I can find a file that does something bad, and I can block it. So, we're focused on malware or file-based protection, and most of our traditional models don't address file-less attacks. Traditional models look at an individual file. Well, what happens when I get a Word document, or a PowerPoint document, or an Excel document, that has built-in macros or scripts? All of us use files that have macros and scripts in them. The file itself isn't bad, but when the script runs, or the script executes, it's a file-less-based attack, and so most of our traditional models don't even look at those particular aspects.

In-memory attacks, that's another area where it's considered a file-less-based attack. So, all of our traditional models are ignoring a lot of these particular attack vectors, and that's why I think we're seeing more and more machine learning coming to the forefront in a lot of products.

Manek Dubash

Simon, you want to weigh in on that?

Simon Crumplin

Yeah, I'm going to weigh in. So, what surprises me about our industry is, we spend so much time talking about malware, and everything seems to be about catching this infection that can potentially harm us, right? But malware is generally a tool. It's a tool to give access, to disrupt, to do something, right? But actually, to materially breach an organisation, malware's just the start, and we spend all our time and focus around this - I call it threat propaganda, right? Our industry is exceptionally good at going, the world is going to collapse, and you've got to buy this thing to stop this threat.

Well, reality is, not all threats are risks to organisations, and if we think, going back to this thing about intrusion that has materially impacts on my business, if I don't transact credit cards, and I've got a bit of malware stealing credit card data, I don't really care. I've got to deal with it, but it's not a priority. If I have someone living in my organisation because they are living with credentials, so they've become someone that I employ, I have a very, very different problem.

What I see in the industry is this threat propaganda, this lack of context on what risk really is. What is the business risk, and what is the context to these alerts? A finance machine has a very different response to the alert than the receptionist, because they're transacting finance. But what we find in our security industry is, we treat it as a blob. We struggle with, which one do we respond to first?

So, in my experience, we've done some work around AI, and what's AI to us? I struggle with it in meaning. I translate it to automated interpretation, because that means something to me, and actually, when I - when we see it taking it away from this theoretical creativity of malicious software that can bypass very control I've ever got, if we actually look at users, and credentials, and someone being compromised - which is arguably more risky - you start to see the value.

It plays to the UEBA space, or UBA, or however you want to term it, to be able to start to get understanding of behavioural patterns with context, because I know what they do for a job, actually starts to give me a much more reduced alert structure, and something where I know how to prioritise.

But I still come back to this point, where we talk so much about technology - the work we've been doing around getting to a set of evidence and fact for an organisation to determine their risk and their risk appetite with the business is the primary thing organisations have got to do and understand, because they can then make the investments that are meaningful to mitigate those risks.

Jan Guldentops

In that field, legal is quite important. Compliance is quite important. If you look at it - there's a lot of comments on the GDPR, new European privacy legislation, but there are good stuff in there. First of all, you need to report a data leak within 72 hours. Of my customers, maybe 15% is capable of doing that, right? That's where this comes in. That's where it - we're going to change - we have to do a risk analysis. We have to know where our data is, which is another big problem. Technology is only there to help them.

What we see is, as a security officer these days, you really have to be - you really have too much work. Right? Some people, some pundits, say the solution is we need more security people. No, we don't. 1999, we needed more IT guys to solve the Y2K problem. Wasn't necessary, but we need to automate. We need to automate simple things, like configuration management. Like log analysis. Let's not call it [unclear] yet. I mean, if AI - and I call - and we were talking about natural language processing, and machine learning, and maybe, in the near future, deep learning, if that can help me, it makes my life easier and makes me more efficient.

Simon Guldentops

I've got to disagree with the piece - sorry to be annoying - around people. You know, I think in my experience, and we're dealing with large enterprise, people that take security seriously, there is no operational capability or capacity.

Jan Guldentops

It's true.

Simon Crumplin

So, you know, you - let's take policy as a good example. There could be 1,000 policy violations. There's no one to follow them up. So, you're not into overload. We eradicated operations and made everything a project, and as soon as it becomes a project you can't respond.

Jan Guldentops

But putting a nitwit next to it who only clicks next is not going to solve it. We need decent security people. Our solution is pulling people from the street and giving him three weeks of training and calling him a security officer. It's not a solution.

Roark Pollock

But one of the reasons we have so many alerts these days is, we've got a lot better at identifying issues after they happen. There's been such a move to - moving from prevention to detection and response. So, if we're finding things that are going on, we're finding more alerts, we're identifying more issues, now we have more things to prioritise and address. AI is getting better at - or, machine learning, to be more specific. Machine learning is getting better at helping us with the task that people - either would need more people to do, or people just aren't good at. They're monotonous.

We can automate a lot of this with machine learning, and get much better at it, and hopefully, in the long run, maybe start to minimise some of the people that we need, some of the alerts that we're having, to address some of the things on the back end.

Jan Guldentops

It reminds me of the t-shirts. You know the t-shirts? Be quiet, or I'll replace you with a very small script. We can do this with AI. Be quiet or I'll replace you with a very small piece of machine learning.

Rik Turner

The algorithms are taking over the world. I believe there was a question. I think there might be a mic that needs to - ah. There you go. Thank you.

Audience Q&A

Unknown

[Unclear] First, a quick remark. I'm not so sarcastic about the good old days, because I seem to recall having read about it, defence in depth. So, it would be pretty silly to chuck out your prevention and your first line of defence, your perimeters. You didn't mention perimeter, but I recall the good old days of the perimeter. So, it would be rather silly to chuck out the old stuff. But regarding machine learning, at a certain point, if indeed at this point it's popping up all over the place, considering the fact that there are only so many attack paths being used, et cetera, there are so many low-hanging fruits, and all those hackers also are into minimax, yeah, having maximum result with minimum effort, actually I was kind of wondering if you would lean back and talk about machine learning and security, but are in [unclear] the fields, the domains, in active security activities where at this point in time research into the use, effective use, of machine learning would be interesting.

For instance, last night I mentioned it's quite often difficult for a company to find an incident and to translate it into what is actually happening. You have frameworks like [STYX] and [TAXI], et cetera, the kind of frameworks to render it into a logical representation of what's happening. It's very difficult. Lots of people do this, perhaps machine learning can help to see an incident, grasp what is actually happening, translate it into something logical representation…

Rik Turner

Contextualise it.

Unknown

…and distribute it in a threat intelligence environment. For the first place, what kind of applications, what kind of fields of research could actually be of interest today for machine learning to actually help in solving different problems in the field of security?

Rik Turner

Okay, guys. First question, please. You can take that one.

Roark Pollock

There's probably three areas that, from our perspective, we're seeing it used, and I'm coming at this from Ziften's perspective, which is we're all focused on end-point protection, whether it's a traditional laptop, desktop, or server infrastructure, virtual machines, you know? Machine learning is doing a few things. One, it's being used on the prevention side, as we talked about, being able to prevent malware, being able to prevent file-less attacks, and that's really all about machine learning, pattern recognition, being able to not only necessarily replace signatures but go beyond what signatures are doing today.

Signatures work, like I mentioned earlier, but it's the areas where signatures aren't working which is unknown malware, you know, before something has been identified in the wild, or file-less attacks where signatures aren't being used. It's being able to go beyond what the signatures are doing today in the prevention side of things. It's also being used from a detection standpoint. Whether it's behavioural based, like UEBA - or UEBA - or it's just basic detection and response on end-point devices. So, we can use it to detect and block from a prevention standpoint. We can also use machine learning to detect things after the fact. So, post-execution detection that we then have to go respond to because something's already happened in our world.

We were talking this morning about it, you know? I think there's still a window for machine learning or artificial intelligence in the world of, how do I look at the overall posture, or the risk posture, of my organisation, the vulnerability posture of my organisation, and make some intelligent decisions about where do I start? I can't do everything at once. What are the highest-priority items that I need to be addressing based on the criticality of devices I'm looking at, the criticality of information that's residing on those devices, and perhaps the criticality of the vulnerabilities that I have, or the risk that I have, based on that device, how it's configured, how it's not patched, et cetera. I think that's still an opportunity.

Jan Guldentops

Basically, setting priorities is a very good first start, but you can use it almost everywhere. One of the things - and this is a network event, but for instance, code review. Code compliance is a perfect area for machine learning and natural language processing. Stuff like, I need to comply to GDPR, I need to make a register of all our personal data, I don't have a clue where they are, right? You can use natural language processing for that. You can do - every field, you can build a tool with one of the toolsets that are available help you do your work, and we're going to see great applications, but it's going to be something limited. It's not going to be the general big solution. It's going to be one thing being done better by a computer, by artificial intelligence.

Simon Crumplin

It's a feature, isn't it?

Jan Guldentops

Yeah, it's a feature.

Simon Crumplin

There's a feature on an end-point, there's a feature with user behaviour. We haven't quite got that ability to bring it together and put context, predominantly because people's data isn't in as good integrity and good quality.

Jan Guldentops

They don't know where they are.

Simon Crumplin

If you think you - we have this silly concept of - silly, sorry. We have this previous concept of, let's put every single piece of data into a data lake, and then run some really, really smart things over the top of it to pull out the needle in the haystack. It's generally - I don't know any other part of a business that works in that way. If you want to understand your financial position, you'd do an audit. It's evidential. It's fact. If you do an audit in your IT systems, evidential audit, you're starting to get understanding of how you operate.

You were mentioning a second ago about resource or bringing people off the street. The best security people we find are IT operational people, because they understand how the business functions, and they understand - when they're educated - the impact of behaviour. As soon as you get to that position, actually they bring context to the organisation and how to respond. Operational risk undoes a lot of these next generation technologies. It's a fascinating thing that I [unclear] control, and then actually my IT ops completely screwed it, so why did I even buy it?

Or more concerning we find is probably 50% of the tools that people buy and implement are partial. They are not finished. If you think about how we're evolving with SD-WAN and all these wonderful cool things, you've got to keep instrumenting as your business changes. I was having a conversation with the CEO at a law firm that we work with, and we were talking about his vision, you know, and how IT is playing into it. He said, yeah, we've saved so much money in moving to the cloud and moving to a sexy network. He said, we've saved so much money. Right, but did you reinvest that money back into security, because you've completely changed the practice of your organisation, and you’re still securing yourself like we did with the home user internet. A firewall, some [idea] - you know, all in a building. [Unclear] suddenly a fabric.

Jan Guldentops

I think there was a question down here.

Rik Turner

I’ve got a question over there first of all, and then one down there.

Unknown

In fact, it tails really nicely onto that. So, to that point of having to create an architecture in this digital age, and to what you mentioned earlier, Simon, around actually deciding trust, and what you were saying, Jan, around automation - what is the broad opinion on building zero-trust architectures based on automation where you can use artificial intelligence or machine learning to drive a trust engine that determines whether people should be able to access resources?

This is actual, a live example, where there is a - let’s call them a financial company in the US who has their distributed users, who need to access SAAS applications, and their ability to access those SAAS applications is based upon a zero-trust network that determines whether they can access it on their end-point being secure. So, they may be judged on their location, on whether their malware or operating system’s up to date, et cetera, and if they don’t meet that threshold they’re not able to access, so you can reduce the threat vector. What’s your - do you see [unclear] working on that sort of…

Simon Crumplin

That’s been available for years. We keep trying to reinvent technologies, obviously to sell the, I presume. But if you take a step back, role-based access has been in Active Directory since the inception of Active Directory. No one configures it properly. Identity and access management is not a new concept.

Jan Guldentops

The problem is, it’s not a technology problem. It’s a problem of having management accept that they can’t get in, for instance, or that people don’t want to implement your role-based systems, or whatever. It’s work, and they don’t want to do it.

Simon Crumplin

The fundamental problem is, the risk is an IT risk, and it really isn’t. It’s a business risk, and we, for some bizarre reason, we are unable to communicate the impact of that risk to our business, so they become accountable.

Roark Pollock

It’s a people risk, and it...

Simon Crumplin

It is. We did this thing where we audited a company, and…

Unknown

Hang on, there is no risk. The evidence is clear. What we've seen over the last three years is a number of research studies highlighting the fact that security breaches, security vulnerabilities, database theft, credential threat, credit card loss, has zero impact on the bottom line of companies, has zero impact on share price. So, the biggest challenge that security professionals have - and I love listening to security professionals wail on about how important they are - is, you actually don't matter. No one cares about IT security because it has no impact on the business at all. In fact, companies who suffer major breaches, and get substantial amounts of press coverage because of it, actually grow as a result of that business.

So, the real value of artificial intelligence is only one thing, in terms of security. It reduces the cost of your lousy products to make them, so that we don't need idiot professionals, and we don't need expensive consultants. All we need is door locks and closed windows, and that's the value of machine learning in security.

Jan Guldentops

I think he's a network guy, so we can spend it on network consultants.

Unknown

Tell this to [unclear] in the Netherlands. They didn't protect - they didn't protect their servers with antivirus. It got hacked, and the business went out of business.

Roark Pollock

Which company were we talking about?

Unknown

Actually, the Netherlands, the government still uses certificate authority, and that was quite a problem, so [unclear] out of business.

Roark Pollock

Oh, you mean - you mean - yeah.

Unknown

But still, they went out of business.

Jan Guldentops

No, they were sold to a Belgian company before they were breached.

Unknown

Actually, the Belgian company changed its name

Jan Guldentops

Belgian security company by the way.

Roark Pollock

Well, that's certainly an extreme viewpoint, I think, of cyber-security.

Rik Turner

They're all the megacompanies.

Jan Guldentops

There is a point in there. It's extremely put, but there is a point in there. I mean, companies can still survive quite a lot of breaches.

Roark Pollock

But I think, if you draw it to a physical - let's compare it to a physical part of the world. If I'm a financial institution, banks don't like to get robbed, but we still pay for securing our banking institutions so that they don't get robbed, right? Banks get robbed all the time. They don't go out of business. I think it's perfectly analogous to cyber-security as well. You don't want to get attacked, you don't want to get robbed, you don't want to lose all of the information that you're collecting, whether it's intellectual property that makes you less competitive, or it's your customers', your partners', or your employees' data. It's the right thing to do, whether it's putting you out of business or not.

Jan Guldentops

Plus, there's the question of trust.

Unknown

It doesn't really matter if you lose the information because there's no consequences. All you have to do is have enough security to stand up in a court of law and say, I did lock the door, your honour, it's not my fault. That's the security - this is where security needs to change. It needs to say, I need to do the minimum level. I need to reduce the cost to a point where it's a good-quality door lock, and then I'm done.

That to me is the secret of machine learning. All I have to do is put a system in place with some moderately competent technology capabilities, so I can stand up in a court of law and say, I put stuff in place. Yes, your honour, I locked the door and closed the windows. It's not my fault, pay my insurance policy.

Roark Pollock

So, I think you're right on in two points. One, you said earlier, machine learning helps reduce the cost of our overall ecosystem. Absolutely. There are some areas where it absolutely reduces the cost because it works well where - in areas where it's very manpower-intensive. So, if we can use machine learning to help reduce the overall cost, every industry in the world does that, right? We want to reduce the overall cost of what we're doing. So, I think you're right on on that point for sure.

Then, the other one is, what is the right amount of security? There's on 100% security. No matter how much money and how many tools we throw at it, there's no such thing as perfect security. So, what is the right amount is a great suggestion.

Jan Guldentops

Just to fill into that, if you look at the new privacy law, European privacy law, how much is spent on explaining how you should secure your infrastructure and your personal data? Guess. It’s less than one page, A4. It says you have to have structural and - I don’t know, but it basically says you have structural and security measures up to what’s affordable and what’s - sorry? And - yeah, the state-of-the-art at that moment.

So, there’s a big point in that. But the problem we’re facing is - and that’s one of the counter-arguments, is saying it doesn’t matter, is the question of trust. That is the biggest issue. Trust of the customer. It’s something completely uncontrollable and unpredictable. Maybe we should put machine learning on that. If you look at how easy a large entity like Facebook, for instance, in Europe, can lose the trust of its customers, or not, because for the same thing, nobody cared, right? That’s the only big argument, I guess.

Oliver Schonschek, Insider Research

Oliver Schonschek from Insider Research. Maybe just coming back to the comparison of signature-based security and machine learning-based security, or the combination of both, first point, wouldn’t you also agree that machine learning looks for patterns. So, something else like signatures. So, they build patterns not of code as the signatures are, but from behaviour or anything else, different factors. So, one idea of course would be - and you had this point - that artificial intelligence or, now, machine learning, can help building signatures faster, signatures of whatever. They could do code analysis and help again the solutions based on signatures to become faster, so not a gap of two weeks or whatever. So, they will make the detection faster.

Maybe you can also tell us a bit about the false-positives, which are quite high with machine learning-based solutions. So, it takes time to find the signatures, but you have not so much false-positives, and you lose time with artificial intelligence-based solutions because it makes a lot of mistakes, at the beginning at least.

Unknown

[Unclear]. Sorry. You have a percentage to tell us about the using of machine learnings for the false positives, because it’s a problem for everybody. For a manager, for a security manager, it’s the false positives. So, does that help, or not? Well? Do you have some figures to tell us, or some numbers?

Roark Pollock

Not off the top of my head, no. I think, from a prevention standpoint, machine learning is very effective, very, very low false positive rates, as you’re using them as an antivirus replacement, effectively. Where machine learning like any other detection capability, whether it’s behavioural-based or not, can produce false positives is when you’re looking at detection. Post-execution detection. So, if you’re using it as doing post-execution detection, it’s always a line of, where am I on the confidence scale of this particular alert on whether it’s malicious or not, and how much false positive - yeah, how much false positive can I deal with?

Do I want to cast a very big net and catch everything that looks even remotely suspicious, or do I want to be very confident in the things that I'm catching so that I'm not overburdening my IT staff and creating alerts that I have to follow up with that might be unnecessary?

Unknown

Yeah, just to pick up on points with regards to whether you think security is relevant or not, the key thing - Jan mentioned earlier about - you mentioned the word tools, right? The important thing is that security admin guys shouldn't basically sit behind what they buy and think that that's the end of the job. That's just part of the process, right?

So, machine learning to its Nth degree could never be more than 99.999% accurate, and in the case of machine learning on end-points, it's always just going to be a best effort, which is fine because that's still better than nothing, as far as I'm concerned. But obviously, humans aren't rational in their behaviour from day to day, and neither are Windows 10 laptops.

Jan Guldentops

The whole security industry is going to be a best effort, always.

Unknown

Absolutely, yeah. Just like anything else.

Jan Guldentops

One of the big evolutions between the '90s and now is that the security industry has become more modest. In the '90s, you could have some arrogant pricks on a stage saying they had the solution for everything, and…

Unknown

I went to Alan Solon's house in 1988 and he claimed he could stop all viruses, right? He didn't even invite me for dinner afterwards.

Jan Guldentops

That's the difference. We got modest. We know - I mean, we - the industry has got modest, and we know what we can't do. That's the important thing. Artificial intelligence is a tool. If you use it correctly and you implement it correctly, it can help you. If you use it badly, it's going to make your life worse, right?

Unknown

Absolutely, yeah. It's like...

Jan Guldentops

It's not magic.

Unknown

It only involves one human to basically screw the whole thing. So, for example, computer-controlled cars, in theory, would work perfectly if there were no humans driving the cars as well, but it only takes on human driving a car and the whole thing's screwed, right?

Unknown

…thinking about the car, it's not anymore machine learning, but much more deep learning. Very deep.

Unknown

Particularly if you analyse Portuguese driving, yeah.

Unknown

Yes, but what about at the end? You will always need a human behind the machines, because machine learning can help you, can reduce the [unclear] attack, the false positive, but at the end you need a human to decide in which direction on what to do. We don't have any AI, good AI, to help you to do that. It's just the human being.

Simon Crumplin

I agree. The elemental view where you get alerts coming from an end-point, or this, or that, you have to correlate them together to give you some form of meaning of behaviour, and actually, what else happened? We did a piece of work where we tested 20-odd next-gen end-points, and it was really interesting. Yes, they stopped malware to a certain or lesser degree, but in a targeted intrusion, you know, I've got a human driving into my business with access and is now doing stuff, when we went to each vendor and said show me what we did, we started at three o'clock on a Friday, they all came back 60% to 75% correct in what we'd done, but 25% missing. It was fascinating, right? With all their machine learning, and AI, and super brains, right?

So, they're asking me as the CEO in this instance to make a risk-based decision on what I do next. What you learn is, you have to augment it. So, Windows has great efficacy, but - event logs, and security events - but there's only 29 of them, not the wave of them that you could consume. So, you become very specific on what you wish to ask to verify whether you've got what you thought you've got. Did McAfee, or whatever - it doesn't matter what vendor - miss something? Because you're trying to get a wraparound, the picture to say yes, it was an alert, no other indications, done. Don't need to worry.

So, it's that correlation and that assurance that I think starts to get to some sort of position of understanding. I also think we - I keep banging on about this. Posture and hygiene, and operational risk, which does play to the gentleman's position over there around security is not necessary - if you can't demonstrate good security posture and hygiene, stop putting snake oil on, because it doesn't do anything. You've got to get control of your estate, and know what's supposed to be running, you know? Otherwise you just - the last few incidents I've dealt with, there was no malware. There was none. So, they compromised credentials and lived off the land, that great operational thing that said, yes, we used a remote access tool called BNC a few years ago. We just didn't clean it up. Well, guess what? They lived in that.

They went around the estate completely undetected and stole data. It was quite interesting, because it did impact this organisation. It was a legal firm, and when those things become public, they - no. So, it does happen. But reputation is a lot in certain industries, and I think it also plays back to this thing of who the victim is, coming back to your point. Financial fraud, it's not new crimes. It's just a different medium to access these things.

In financial fraud - you know, I had identity fraud done on myself. Once I got the money back I was no longer the victim. It was the insurance. It was managed by the banks, which we then pay for our charges. So, again, I think it's correlation and getting control, because security isn't as complicated as we make it. It is relatively simple once you get to a set of data that an organisation will engage with.

Jan Guldentops

Yeah, and especially the KISS principle…

Simon Crumplin

Keep it simple, simplification.

Jan Guldentops

[Unclear] is - applies to security. I mean, I do a lot of what they call tiger team [unclear], and I'm usually sat in for the social engineering bit, and usually I win from the technical team. Why? Because you walk in somewhere, you're friendly to the people, and they do everything for you. Right? They give you everything, they do everything. It's so simple to counter that. It's give people procedures. But you don't. The only procedure they have is, oh, she did that. We have to fire her. But nobody told her she couldn't talk to me or give me data.

Roark Pollock

Yeah, security is hard work. There's no doubt. There's no product out there that's a silver bullet. There's no machine learning capability out there that's a silver bullet. It does help us - to change your analogy a little bit - to find a needle in a needlestack, which is a lot - you know…

Rik Turner

Rather than a haystack.

Roark Pollock

Rather than a haystack, it's like looking for a needle in a needlestack, which is what security is often like. It does help us with those types of activities, but it's not a silver bullet, and there is no silver bullet technology out there.

Simon Crumplin

I've had a couple of conversations with organisations that have bought one of everything. They've got 60 technologies covering everything they can cover to try and get control, and to catch the thing that's material. Not the commodity bit of malware from a botnet that they can eradicate because it's just an infection. In all of those conversations that I'm having a year later, they're all going, I can't do it. I don't have the resource, I'm not a software development company, to try and bring these 60 tools, and these 60 islands of people and knowledge, together. I need to get it down to six or seven technologies and leverage that integration.

Because I think it feeds the - how do you get benefit out of machine learning and AI, if I'm - if I have confidence in the integrity of my data, I can start to make predictions. If I've got 60 different sources I never have confidence, because it's too broad.

Roark Pollock

Well, it comes back to the point the gentleman made earlier, is what am I trying to achieve, and what is the right level of investment from a tool standpoint, from a money standpoint, from a resources standpoint, to achieve that end-game? You can't just spend everything a company makes on security. You've got to be able to optimise it and figure it out.

Rik Turner

Folks, I think I'm seeing that we're coming to a natural end of this conversation, so I would invite you to come coffee. Thanks very much to the panel. Thank you very much to the panel. Thanks to Rik in particular. Let me just quickly, before we shoot off for coffee, remind you of the Twitter handle, #NetEvents18. Get tweeting, and the winner of the most tweets will win a lovely cork laptop case. Coffee is downstairs, I believe. So, yeah. See you later.

[end]