

Speaker 1:

Throughout this summit so far, we've defined the threat landscape, unpacked collaboration efforts, and talked through some of the strategies that can be used to strengthen cybersecurity.

This afternoon, we're going to look ahead to the application of those strategies, future possibilities in the world of cybersecurity. To kick off our look ahead, we'll be exploring what AI really means for the future of cybersecurity, important lessons from the international threat landscape, and more.

Before I welcome our speaker up to the stage, we have a polling question for you all. Will the recent advancements in artificial intelligence be revolutionary or evolutionary? This is... Well, I'm not even going to bother unpacking the question. I'll let you all answer how you think, and then Jeff can opine on it. We can explain more about what we mean here. While you all are starting to think about this, I'm going to welcome up Jeff Greene, Senior Director for Cybersecurity Programs at the Aspen Institute. Jeff, come join us.

Since people are jumping in and answering the question we have up here, this one actually came from you. Explain to us what you mean here when you say revolutionary versus evolutionary.

Jeff Greene:

I'm watching the poll changes as we go. Oh, here we go. When we were talking earlier, I said, "I was promised the blockchain was going to fix the world." Still waiting for that to happen. I count myself as a little bit of a skeptic, as to how fundamentally revolutionary generative AI is going to be. Not because it's not an amazing technology that doesn't do amazing things, but I'm waiting for a good explanation of how it will do new and different things, as opposed to do the things that we're doing right now differently, more efficiently, maybe more dangerously, maybe more safely.

I don't have a good answer to this question. It would probably depend upon who I spoke to last how I'd answer it, but I think it's something that we need to bake into our conversations, is what is the change? I've heard enough smart people who I trust saying, "this changes everything," that I won't be surprised, but still waiting to hear.

Speaker 1:

Yeah. Well, so let's get into it more and talk about artificial intelligence because I know you and I both have pretty strong opinions on what is, and what isn't, and the state of it today. Before we get into that though, for those out there who aren't familiar with you, can you give us a rundown on your bonafide? Where do you come from, what are you doing in these days?

Jeff Greene:

Sure. Off the record, I was a lawyer at one point in my career. I started as a lawyer and got into government, actually back during the Katrina investigation on the hill, but ended up doing cybersecurity. For those of you who've been around a bit during the Lieberman, Collins days, I was in the Senate Homeland Security Committee working on that legislation. Spent about eight years at Symantec running their global government affairs and policy, and then was over at NIST at the National Cybersecurity Center of Excellence. Then in March of 2021, was detailed to the White House where I ran the defensive policy work out of the cyber director there working for Anne Neuberger. That's where I was until last summer. Very slow time in defensive cyber work in 2021 and 2022.

Speaker 1:

Excellent. As we chat here too, we're getting interactive with the poll. I want to get more interactive than that with all of you. Feel free, if anybody has questions out there, if it's more of a comment than a question, maybe keep it brief, but feel free to raise your hands as we go along. I'll be looking for you. If I miss you, start waving around aggressively, but we'll keep this interactive if anybody wants to join in the conversation with us.

As we're talking about what AI is and what it isn't, one of the things, the way you were talking about it at first, is it revolutionary or evolutionary? Because so many things, I hate the term artificial intelligence because these things to me are not intelligent. I like to call them in initial caps, very good algorithms. They're doing really important... Well, not necessarily important all the time, but they're doing really interesting things and difficult things. They're making huge leaps, but they're not intelligent. Personally, I feel like that can lead people astray in thinking that they can do things that they can't.

You're in and around this stuff. You've seen how it's being used for cybersecurity and the like. What can it do today? What do you think it's going to be able to do next year?

Jeff Greene:

Leaning a little bit back on my Symantec days, we had a lot of internal debates on what to call AI, versus ML. But thinking through that lens, what I can see it doing soon, if not now, is making it easier for humans to keep us safe. By that, I mean if you think back to some of the major breaches we had in the middle of the last century, many of them were caught by security tools, but the security teams were overwhelmed with alerts and weren't able to pick out the wheat from the chaff.

Ideally, you have systems that are able to make it easier for humans to interact, engage where they need to, maybe in some situations to be able... And when you think about endpoint detection and response, on the response part maybe have some more automated quicker responses because humans are never going to be as quick. I think we are probably seeing that and will soon. I don't know what the next generation of that is.

On the criminal side, what worries me immediately is something as simple as a phishing email, the quality of the phishing emails. I'm doing a workshop in the coming weeks for some folks very new to cybersecurity, and one of the tools that I was using talks about, "Oh, look for bad grammar, look for people who may not be native English speakers." That will be less of a problem. And not only that, if you want a tailored phishing email, write one, you write a phishing email for someone who's worked in the White House and is interested in hockey, put those two together. You might get something that's more likely to hook me than it would be someone else.

Speaker 1:

It'll be a whole new generation of script kiddies.

Jeff Greene:

Yes.

Speaker 1:

Because not even just the phishing, but even writing code, you could write malware without having to know a line of code.

Jeff Greene:

Yeah, well and I've heard various smart people on both sides say, "This stuff is writing good code and will get better," and, "This stuff is not writing good code." I don't know which it is-

Speaker 1:

But that sounds like human programmers too.

Jeff Greene:

Yeah.

Speaker 1:

We're talking about the evolutionary side. If this is an iterative progress, for AI to be useful for the defenders, what does it need to be able to do? You talked a little bit about using it in terms of helping cut through some of the noise, and get to the signal, and all that. A couple of buzzwords I like to use for that thing is the difference between automation versus orchestration. The idea of you can automate the whole thing so a human doesn't have to touch it, or do you orchestrate it so that it's working with the human and they're working in concert together? Are we anywhere near full automation for cybersecurity? Can we just set it and forget it?

Jeff Greene:

I would challenge whether we will ever reach full automation. I was thinking when you were talking, that it's a little bit of both. If we can get to the automation where many... I hesitate to use the word attacks, but it's many attacks, intrusions, whatever you call them, we're able to have an automated response because it is of a kind we've seen before. And most compromises today are still based on known patchable vulnerabilities, old techniques, known malware. So if we can get to a point where many of the responses can be automated, it will shrink where we need to have the orchestrated piece.

Speaker 1:

In some ways... This is where, again I go back to the very good algorithm thing, where we can divorce the buzzword of AI and this new thing coming on board with what's actually happening in the real world. A lot of that is automated in many ways already, isn't it?

Jeff Greene:

Yes.

Speaker 1:

If I go to sign on to something and I sign in incorrectly three times and it locks me out, that's an automated system. There's no human doing that. How is AI going to bring that to another level? How do we create algorithms that could actually free up human time or do things that are better than humans? Where in that line does that get inserted?

Jeff Greene:

When you are logging in, ideally the system, if it's a good security system, is checking more than did you get the password right? It is looking, is it typical for you to log in at that time from this location? But those are the very obvious data points they can examine as you get deeper into all of the characteristic. It's the thing that creeps us out about the targeted emails we get because they know us so well that we

can be targeted for advertising. If you have a system that knows your habits so well, it may pick up on an anomaly in why or how someone's logging in as you that no human would get.

Even more important than you doing that, it is monitoring the entire system to see is this system acting in a way that is just enough out of the norm? That's where the human gets in. At some point, you need to set those rules about what counts as out of the norm or not. Although ultimately, ideally, I guess you have the AI systems figuring out where to put those boundaries-

Speaker 1:

Over time, yeah.

Jeff Greene:

It is looking for anomalies across massive data sets that no human will get. Now, I don't have a good answer if that's ML or AI, and when it becomes one versus the other.

Speaker 1:

Let's talk about that distinction while we're parsing buzzwords. In your estimation, what's the difference between AI writ large and machine learning?

Jeff Greene:

Boy, I'm going to go out on limbs here. I wish I could remember what I was taught to say at Symantec. ML is more rule following, following a set of procedures. AI is drawing conclusions and coming up with something new and novel maybe that was beyond the rules. Again, I don't know if these are true, but you read, I can't remember which LLM, but it had learned a language that it was never told to learn.

Speaker 1:

Yeah, that was a bit on the bogus side because it didn't really learn the language so much as it was able to find and pull threads from. It's a parrot machine versus learning. This gets to what we were talking about before, is it intelligent if it can't actually learn?

Jeff Greene:

Yeah. Well we talked about the phrase using hallucinations for bad answers and that putting a human construct on top of something that is still just ultimately a set of rules building on rules. We are at the point where they're so complex, that no human can ever unwind how we got to the point of making that decision. I'm sure there was a logic to it within the logic of the MLR AI system.

Speaker 1:

Sure, but even with humans though, it's one of the things that makes humans better than computers right now, is that you could have those higher order processing where I don't know why I make certain decisions. Given a day, given the time, I could probably break it down, but...

I want to bring it back to the practical side of all this because my goal here isn't to just dump on AI for the whole session, but to get to where these advancements in technology can change things. You mentioned blockchain before it and previously I've heard you talk about this in an interesting way. Is artificial intelligence where it's going now, these advanced algorithms and what they're going to do for

cybersecurity, is it blockchain like you said before or is it, I've heard you compare it to zero trust? Can you break down the thought process and what you're talking about there?

Jeff Greene:

Yeah, I think what we're talking about there is when I think of blockchain, blockchain technology is everywhere and doing lots of good things, but you go to RSA or other cyber or tech events, you hear about that a lot less. The zero trust concept, which has been around for a while, but has really taken over, there was zero trust strategy that the federal CISA put out. As a concept, I think of that as very enduring, whether you call it zero trust or not, this idea that you never trust an interaction between a human and a device or between a device and a device-

Speaker 1:

That's entering your security posture on identity management.

Jeff Greene:

Identity management or access management.

Speaker 1:

Access, that's better. Yeah.

Jeff Greene:

Anytime the system's trying to do something, you assume it is, if not malicious, not trusted, and you're verifying, and you're using that type of data we talked about to kick it out maybe to a human. That's where I think of AI as having, whether it's AI, ML, or the hybrid, a very long-term positive effect on that side because it will allow us to make zero trust even more effective. I think the phrase is so popular that we're going to have enough TV shows in the next five years about it that it'll never go away.

Speaker 1:

It's going to be the next CSI?

Jeff Greene:

Yeah.

Speaker 1:

Zero Trust Miami. Let's talk about some of the advancements we've seen. I think critical infrastructure's a great place to look at these because it's something that resonates so much with people and we've seen so many real world things. We've got Colonial Pipeline, we've got the... Remember the Bowman Avenue Dam in New York, where they got the wrong dam, and there's just a slough and not an actual dam. What could have been?

Jeff Greene:

Yeah.

Speaker 1:

All the way to Ukraine years ago, when the cyber attack knocked out the power grid. When we talk about cybersecurity of these critical infrastructures, this is another place. Can artificial intelligence improve this or is it going to actually make things worse? Not just from the attacker side, but just the confusion among defenders.

Jeff Greene:

Yeah. It will definitely facilitate attacks. It will help on the defense side when organizations are using it. And by that, I mean if you think back even long before, if you go back to maybe 2010 or so, there was an intrusion into a water system, South Houston water. That was one where using a common search engine, called Shodan, a gentleman in Eastern Europe was able to find it, determine what operating system it was using, Googled the type of SCADA system it had, and found the default username and password, which had not been changed, and logged in, and had full control over South Houston Water District. Now, this person was doing it to prove a point that it could be done, but AI is not going to help South Houston because unless they're deploying new technology-

Speaker 1:

Or if they deploy an AI system that would go around changing and admitting guest passwords, that might be something.

Jeff Greene:

Yeah, I mean if you think back to the Mirai botnet, which took over IOT devices in I think 2016, it was using default username and passwords. If you had a device where the default password was one, two, three, four, you would've been more secure if you changed that to password. That's the type of simple things we're still worrying about. I don't mean to pick on Colonial, but again, they were using single, or the Philadelphia Enquirer, single factor authentication with probably information about sources that is supposed to be confidential.

If true, that is a shocking level of, I don't know if it's incompetence or it can't be because it's just lack of attention to it. Someone chose not to invest in that basic security. AI is not going to help them if they're still using single factor authentication, but it will make it easier for the attackers to try to spearfish, or write the malware, or other pieces.

Speaker 1:

It sounds like from your perspective, from the attacker side, AI's going to be used for natural language phishing attacks, maybe some script kiddies, like I was suggesting. From the defender side, it's mostly anomaly detection. Are there any other good uses for it?

Jeff Greene:

The response too. I think we have to assume that at some point if we are able to use it to automate response technology, we have to assume that attackers will be able to use it to automate when it sees it's being stopped in a certain way to be more effective.

Speaker 1:

The dynamic.

Jeff Greene:

This transcript was exported on May 22, 2023 - view latest version [here](#).

The attack hits you see that are, whether it's ransomware as a service or the other types of attack tools you can buy, will have much more effective ways of evading security. It is the constant cat and mouse game.

Speaker 1:

Is it then just about having better AI than the attacker and just letting that be a race between the two?

Jeff Greene:

For those who are deploying the best technology, yes.

Speaker 1:

We're talking about the federal government, which depending on the agency and the program, is either 15-year-old technology or best in breed. You have to span that.

Jeff Greene:

You'll always be left behind the organizations that don't have the funding, the capacity, or the intent to put in the best-

Speaker 1:

You brought up the Enquirer and I have no special knowledge. I've never worked there. I don't know why they had single factor versus other things, but most often when you have a security failure in an organization, it's because they don't see the return on investment in investing in it.

Jeff Greene:

Yeah.

Speaker 1:

If you're an agency out there and you want to invest in artificial intelligence products, the brand name stuff, for cybersecurity how do you go about doing that in a way that you can show a return on investment, that when you're doing your congressional budget cycle, you can say, "Look, we spent this money, here's why, and here's the effect it's going to have," and not just lead to budget cuts, and them thinking you're wasting your money on shiny objects?

Jeff Greene:

It's how do you prove the negative of you've prevented attacks?

Speaker 1:

Especially for an untested technology?

Jeff Greene:

Yeah, I think probably if I were giving advice to those, tie it to the current strategies that the administration, that OMB are putting out, show how you are achieving whether it is the... I'm blanking on the acronym, but some of the benchmarks that NIST, whether it's 853 or something else, show how you're achieving the security ends that you are required to meet, whether more efficiently or otherwise.

I think also, point to if you see attackers moving into this space, point to the type of intrusions, attacks, techniques that you will now be able to defend. And last thing actually is, if you can show how this will make your current workforce more efficient in detecting and spending their time on the really high risk incidents.

Speaker 1:

If you can show direct cost savings.

Jeff Greene:

Yes, cost savings or were able to focus on that wheat from chaff line again, our people are spending their time on the most important things.

Speaker 1:

We've got about five minutes left. I was just going to see if anybody had questions. We got one right here.

Speaker 3:

So-

Speaker 1:

Oh, we got a mic for you.

Speaker 3:

Can you hear me? Looking at AI with respect to tools and technology, as somebody mentioned earlier, those of us... I'm a computer science guy, the source of training for the AI model is just as important as what it does.

Jeff Greene:

Yeah.

Speaker 3:

Is there a strategy to vet the AI that some of these security tools are bringing to the table? Couldn't we, just as easily let a bad AI into our ecosystem if we don't have a strategic way of vetting where that AI got its source of knowledge?

Jeff Greene:

That is an incredibly important point. Yes. What is the data, data poisoning, et cetera. I don't know the answer, but if you look at the NIST trustworthy AI, I think it's called the Framework, I believe they touched on some of that.

When I left the NCCoE or when I got to the NCCoE, we were just about to stand up a trustworthy AI lab that's in their facility in Gaithersburg. Then COVID hit a month after I started. And last I heard, they were putting it together, but those are essential questions. Without answering your question, that is an-

Speaker 3:



Are we going to [inaudible 00:19:18]?

Jeff Greene:

That's the visibility piece that I think you heard some of the members talking about after the hearing yesterday, the day before. How do we understand what the AI is using and built on? How can we trust it? And I don't think there's a good answer right now. I was at a global event a few weeks back and everyone around the world is asking essentially these same questions and no one has a great answer yet.

Speaker 3:

Everybody's selling us tools with AI built into them-

Jeff Greene:

Yeah. And governments are ready to regulate and move out on it when they're not sure what it's going to look like.

Speaker 1:

It's also they're selling it with AI built into it. One of the things is, if you're a computer science, if you're into technology, if you're a CTO, you probably have a good idea of what AI actually is and isn't. But the average contracting officer, do they really, when somebody is trying to sell them something that says it has AI in it, how do you even understand a thing if you're not a high-end technologist who can break down all this stuff? How do you not just get built on buying stuff?

Jeff Greene:

I don't mean to impugn contracting officers because they do yeoman's work, but most of them end up reading a lot of materials that are written by the marketing parts, so you have to rely on your IT, your security team to tell you what they're using and relying on.

Speaker 1:

And it's a lot of what we do at Nextgov too. I know a lot of our audience is the contract community, and one of the things we try to do is bust some of these myths to just help them get through all that.

Jeff Greene:

Yeah.

Speaker 1:

Yeah.

Jeff Greene:

I would be hesitant to buy anything that I'm told is truly revolutionary, fundamentally game-changing, does something unlike what anyone else can do because there certainly are things that are better. There's Lindt chocolate versus Hershey's chocolate maybe.

Speaker 4:

Going with the chocolate metaphor here, the desserts were great. As the co-author or really main mover of EO 14028, you did a lot to really ensure that there were fundamental understandings of what zero trust was, cloud security. Two days ago, as was briefed in front of Congress, a lot of AI proponents asked for guardrails. If you could write or recommend to the government what those guardrails would look like, what would that include?

Jeff Greene:

I am glad we weren't writing 14028 with AI in mind because I don't have a great answer for that. I think the biggest thing, and I need to do a shout-out to the team that really focused on and wrote... It was largely written when I landed there. There was some great people killing themselves on that, who are still serving in government.

To me, the thing I would've said, which government and policy makers often don't like, is move with caution. Take some time before you try to set a hard guardrail. I had a big cyber meeting through Aspen Global Group in November, and we barely talked about AI. That wasn't because we were ignoring it, because this just wasn't the focus of the world then. Things are evolving so quickly, I would take time before you set a hard guardrail and make sure whatever you do has flexibility built in, because the freshness date on any AI guidance is going to come up really quickly.

Speaker 1:

Well, one of the things that came up in that hearing too, was the idea of a regulatory agency. I don't know that we need a new regulatory agency for AI. If government were to regulate how this is used and not just how it's developed, but how it's sold, how it's marketed, all those things, who would do that? FTC does some of that already.

Jeff Greene:

FTC will do a piece of it.

Speaker 1:

You should read the FTC blog, that science and tech blog by the way. They are on fire lately. They've been great.

Jeff Greene:

My immediate reaction is a new structure rarely solves a problem. God love the members who were talking about it, but I would be hesitant to do a reorg to address something, the last question we don't fully understand yet.

The only way I think to regulate it government-wide, would be through procurement. The most recent executive order had some guidance on secure AI. I think the government's still trying to figure out what the best... CISA certainly will have a piece for critical infrastructure, but it depends on what you're talking about, Cyber AI versus how AI's going to-

Speaker 1:

Robotic process automation.

Jeff Greene:

Yeah, I mean the labor department's going to have a very different view on this than the Education Department, than CISA, than DOD.

Speaker 1:

Sure. Yeah. We got time for a quick one if anybody has any more questions. No? Let's just end it up then with something more, the open-ended of this. Is this something to put all the eggs into your basket? If you're looking at trying to revolutionize your cybersecurity posture, you're at an agency, you're at one of the big name, one of the 24 agencies just announced that they got hit pretty hard in a breach that lost a bunch of information about their employees, you want to go ahead and improve your cybersecurity posture, should you be focused on AI or is that just something that's going to come along with the right tools? If you're buying the right stuff, it may or may not have it, or should you focus on that and use that as your linchpin of your program?

Jeff Greene:

Last point, yes. I think you should assume many of the new modern tools that you're going to acquire will have this built in. Obviously, you want to ask along the way. But to the broader question, it really depends upon the security posture of the agency. SolarWinds showed that there were different agencies at different levels of maturity. You unwind back to OPM-

Speaker 1:

And that was still a person who made the determination that caught it too.

Jeff Greene:

Yeah.

Speaker 1:

And did some great reporting recently on that.

Jeff Greene:

Yes, exactly. And if you go back to OPM though, OPM's security wasn't the way it was because no one cared. In part, if you read into a lot of the reports that came out before that, they were not often given the funding to put in the type of security that they knew they needed.

If you are still struggling with MFA, you have to look at whether your systems are capable. One of the big pushes in 14028 and around then, was MFA and encryption. A lot of the systems weren't capable of taking it. Trying to force an agency to put MFA or encryption into a system that just was too antiquated, you need to reset to the baseline.

Krista Rush, the federal CISA, in my estimation, did a really nice job making sure we focused in the White House and in the NSC side of not just saying, thou must do X. The first question is, are thou capable of doing X? That's where I would start.

Speaker 1:

That is a great place to start. You made me think of one other thought in all of this is, you mentioned before if you're trying to sell your program to say Congress for budgeting or your department heads and

This transcript was exported on May 22, 2023 - view latest version [here](#).

the like, maybe having little AI in there isn't so bad because of the direction from the White House and the like, where they do want to see more of that stuff.

Jeff Greene:

The shiny object.

Speaker 1:

Yeah. Get that shiny object, but make sure you know what the shiny object-

Jeff Greene:

Yeah, what it's going to do and not do.

Speaker 1:

Yeah, exactly. Jeff, thank you so much for joining us today.

Jeff Greene:

Yep.

Speaker 1:

We're going to hand it off to the next group.

Jeff Greene:

Great.

Speaker 1:

Thanks, everybody.