MIT CSAIL Alliances | Ellis_Herzog_Project_4(1)

Welcome to MIT's Computer Science and Artificial Intelligence Labs Alliances podcast. I'm Kara Miller.

[MUSIC PLAYING]

On today's show, two cyber security leaders, Amy Herzog at Amazon and Andy Ellis, former chief security officer at Akamai talk about how corporate leaders should think about cybersecurity right now. From the role that AI is going to play--

How do we identify and resolve issues? How do we prioritize the work that's in our queues? This technology is really good at amplifying human abilities.

--to the one vulnerability we should worry most about, and it's probably not what you think.

We run basically the world on email, and it's the most unsafe protocol out there. You keep coming back to phishing. Very real threat. But why does phishing work? Because I could send email and claim to be you, and your mail browser will tell you you just sent mail to yourself.

That's all coming up on today's show.

[MUSIC PLAYING]

The AI transition feels like it's changing everything about the corporate world. But, says Amy Herzog, chief information security officer for Ads & Devices at Amazon, everything may be overstating the case.

One of the most helpful things I have found in the last year is to remind myself that from a security perspective, as a security partner with the business, a lot of the work is the same. We start with the customer the context of the products, the goals of what we're trying to do. We get clarity on what we're doing and why because that all helps you understand what you need to do from a security perspective to make sure that you're successful.

Figuring out what the threat model is, how to build components securely from the start-- the tactics are different, but the strategies are still the same and familiar.

Still, notes Andy Ellis, operating partner at YL Ventures and former CSO at Akamai, AI is about to alter cyber in ways that could be awfully helpful.

One of the biggest challenges I often see in almost every workforce is actually writing. We publish documentation, and it's written by developers who English might or might not be their first language, even if they grew up in an English speaking country, C might be their first language. And being able to write coherent and readable documentation has been a huge challenge all through the supply chain.

Now employees can become editors rather than spend huge chunks of time writing. And Ellis and Herzog see this as emblematic of a larger corporate shift, a shift they think is very promising.

So many businesses have already pivoted to almost a complete outsourced model for much of their enterprises. You're using SaaS. You're using cloud services. You're using third parties. And the way that AI can enable every one of those to operate even if it's only 5% or 10% more efficiently like the way in which that's going to affect our whole ecosystem and create this productivity boom for companies to me is really exciting.

And I think as Amy said, you still have the same challenges. They might just operate at a different scale. If you were outsourcing your-- write your documentation to a third party in Indonesia, or you're outsourcing it to ChatGPT, you still have the same data leak problems. You might just be using more of ChatGPT because of how fast it is.

So let me follow up on that. One of the concerns that I've heard from companies-- I think some scholars think it's overblown, I'm sure some don't-- is even letting people in an organization use big foundation models.

Because they're often not internal. You're putting something on top of something else. And there's this concern what if your data escapes. It's not fully in-house. From the security perspective, what kind of concern do you hear about that? Do you feel like it's a real concern? And it is indeed really on companies minds.

So I think absolutely companies are thinking about this. And that's why you see people who are sometimes using external, and then more and more people are saying, well, how do I bring that in-house so that it's my model that's learning from my data.

And I think you just have to explore. Like, here's one benefit of using models rather than humans is models can be thrown away at the end? Like, if this is something that doesn't need to learn, and if the third party says, yeah, we take your data. We use it to give you a result and then throw away our learnings, great.

But that's the conversation you have to have as a security person is, are they throwing away your data so nobody else has that included in their model, or are they keeping it and growing their model, in which case, I think you do have a problem.

Yeah, just to build on what Andy was saying, I think in any LLM or any GenAI context, you start by asking the question, where is our data? You ask what happens with my query and my associated data? You ask is the output of these models fit for purpose or accurate enough? And then you're starting from a really reasoned place.

And Amy, I wonder from folks that you talk to if you feel like there are, let's say, sectors of the economy even, I think of health care where people-- sure, AI has come in a flurry and the last year people have gotten really excited-- but are there sectors where things, privacy, really, really matters maybe in ways it doesn't as quite as much in some other areas? And I wonder if companies are more reluctant to jump into new technology because they have those kinds of concerns about privacy and security?

Yeah, in some of my past roles doing transformation work across a lot of different sectors, different companies, different sectors definitely do think about their interaction with their customers, their data, all of that in different ways. So yeah, for sure. Any product, any company is within a really unique context centered on their customer and their customers' needs. And that's how we build products here at Amazon. That's how other companies make decisions about trusting their data.

I think the thing that you can do to make sure that you're adapting to any change, GenAi or not, is to really start by building on a strong foundation. Partner the security teams with the business teams in place through the product development process. That's when you'll be able to ask and answer those questions successfully in the way that allows products to be released the fastest.

Amy, let me stick with you. I know one of the things you've thought about is how companies deal with security when they're working with foundation models that act in unpredictable ways. We're in an era that's very different from the way things used to be, or we're entering an era that's different. And when AI is involved, you can have answers or pathways that feel unexpected, they feel random. Is that a different ball game when it comes to security?

Predictable products and unpredictable products are different from each other, for sure. But the LLM, the model interaction at the heart of any experience is only one part of that experience. And so I think the place that I would like to see the industry go is to really think and develop, OK, we have non-deterministic queries at the heart of a customer experience. What is the rest of that? Or user experience.

Because I think some of the places these models across industries are most relevant are places where humans are doing something that can be learned from already. And then you need to think how does the LLM query fit into this workflow, and what are we building around it to make sure that the ultimate experience is the one that we desire?

And I wonder if you think even on the security side, because there's more AI just being used in general, if there's also going to be more AI used in the security solutions?

Yeah, for sure. Security engineers there are not enough of them, and definitely, part of our work is carrying out tasks that have opportunities from an operational efficiency perspective. How do we identify and resolve issues? How do we prioritize the work that's in our queues? This technology is really good at amplifying human abilities. So understanding what particular step you want to make more efficient, I think there's tons of opportunity.

And you talk about that. I feel like there's that very interesting kernel that Amy identifies there, which is like there are not enough professionals in security. So then what does AI do to this equation?

So what AI is really going to do is remove the wasted overhead. I think Amy just hinted at that one, which is, I think, there's this old story. I love this saw about the engineer or maintenance tech who retires from some facility. And after he's been retired for a year, the boiler breaks down, and nobody knows how to fix it.

Gets called in. Puts a piece of chalk and says, bang right here. And they bang it. And he's like here's your bill for $30,000. And they're like what's that for? And he's like $1 for chalk and $29,999 for knowing where to put the X.

And that's what AI will do for us is you can take these processes. You think about a vulnerability that happens. It comes in. You've got to identify every affected system, figure out what the patch is, figure out who's going to implement things. And right now that's humans do much of that work.

But imagine if instead like a vulnerability comes in, and somebody gets basically told, hey, here's the vulnerability. Here's what it's going to take to fix it. Here's the risks of fixing it quickly versus not fixing it. Do you want to proceed, yes or no? So a human still controls whether or not we're doing something. And maybe there's some step they have to take, but we remove all of the overhead that makes us very inefficient as businesses.

And how good is the technology there? Like to what degree is any of it plagued by hallucinations or just errors or, I mean, other kind of thing that can happen with AI.

Yeah, so I think we have to watch for is where are hallucinations likely to show up? And I think the most common thing is anytime you ask AI to name something is where it seems like it hallucinates more often. Like, tell me what version to upgrade to. If it's not pulling that out of a straight API, it might make up a version for you. Be very careful.

But a lot of times it's just, hey, look across a large data set and find all of the things that match this or look similar or make an inference about who it might be. And then you give it to a human. Imagine if I got a notice that said, hey, we've got a vulnerability. We think Amy could be the one who fixes this, and it's the first time we've ever identified Amy.

So Andy why don't you call Amy and just check. I make that call. Amy's like, oh, yeah, that's my system. And now the system can learn and say, OK, great, when it's this system, it's Amy, and the next time we can now automate that. So I think that's the real challenge for organizations is how do you bring AI in the same way you would bring in a junior employee?

You don't bring in a brand new employee and say, here have root access on all of our systems and make all the changes you feel like. You watch them. You train them. You trust them. And you establish that over time, but there are productivity multiplier.

But it sounds, Amy, like you also may have human as gatekeeper or human as oversight person in that equation.

Yeah. I think the thing that Andy's pointing out here really rightly is that the use of GenAI or LLM technologies in a product that's not a thing that is impossible to control. And it's not a thing that's going to stay the same over time. So as you think about constructing a product with security in mind from the start, think about what are the use cases you're worried about.

Maybe the architecture of the overall product has protections or mitigations built in to focus the human's attention really efficiently. And so it makes sense from a scale perspective to have a human in the loop before the output ever goes somewhere. There are lots of ways that you can use this technology in product development that's not all or nothing.

I wonder-- either of you can take this-- but I wonder when it comes to entry-level employees, I mean, in some ways, the technology that you've talked about can take the place of the entry-level employee. And the bonus is like it never goes to sleep, and it never takes a day off and that sort of thing. So it's always looking for problems in the system.

But to what degree does that take out that entry-level staff so that then maybe you don't have entry-level security professionals climbing up the ladder and knowing the ropes in the way they once did? I don't know if that is a concern for either of you.

I think security is a field that requires calibration and high judgment to such an extent that we would never want to eliminate the best path to gaining that judgment and calibration, if that makes sense. They're helpers. They make us more productive. That's OK. We don't have enough humans right now to actually tackle this problem on a global scale. We could use the help amplifying those abilities, but it's not a replacement.

OK, you still want people who can get to each rung of the ladder so that someday they are the people with the oversight.

Right, you need that person who the AI is going to escalate to, and say, hey, look, I saw a million things today, and 900,000 of them look exactly like what we saw yesterday. So here's what I'm doing. But there's 100,000-- here's the 20 that I don't know. Can you go look at them?

And when you think about that entry-level position today, all too many companies basically have that entry-level person just working on the 900,000 anyway. They aren't actually exercising their judgment because they're following a checklist, and they're doing a thing, and it requires them to go above and beyond their job to get the experience to move to the next level. What I think this will enable is that becomes their job is gaining that experience and learning better judgment quickly.

So let me stick with the idea of how companies are dealing with cyber right now. Obviously, in big companies, there's a lot of attention on cybersecurity, especially well-financed companies. I wonder-- Andy, we can start with you-- if you worry that as the threat grows, there's a bigger and bigger gap between the well-financed companies and the average small or medium company and how they not just think about but honestly, just have the resources to deal with these kinds of threats?

Yeah, I was really worried about that probably five to 10 years ago and I'm much less so now. I have to say that the migration of businesses away from a centralized IT infrastructure to one that's completely cloud-based completely changes the dynamic.

I have my own side business. I have a consulting company. And I'm running everything-- sorry, no offense, Amy-- inside Google. Actually, I'm not. I've got things in AWS too. So I'm actually two-vendor on that one. And my security concerns are, am I configuring correctly my AWS interface and using the right tools there? And am I setting up my Google domain the right way?

But I'm not doing systems administration. That's a huge change from the way we were five or 10 years ago that if you were a small business, like you went and you installed some servers in a closet somewhere, and you never maintain them. I'm not worried about that. Amy's responsible for maintaining a bunch of my infrastructure. I feel a lot safer now.

So really small companies are big companies now.

Right. It's amazing.

It's really incredible. It's such a boon for innovation. And I think from all of my experience in the security industry, it's that lack of attention, the lack of maintenance, the not being on top of patching, it's that stuff that gets you in the end. And how wonderful to have a whole ecosystem now where the small innovator doesn't have to worry about any of that stuff.

Andy, when you sit down with company leadership. If they say, what is our greatest vulnerability? Is it our people, the human engineering piece of it? Is it the software we're running? If you had to identify where you think that biggest gap is, what do you think?

So I think the biggest gap in most companies it's never the human, although, people often start that way. I'm a huge fan of Professor Nancy Levison's work in system safety. And my favorite way to quote her is she says, "human error is a symptom of a system in need of redesign."

OK. All right.

I love that.

And that's where most of our problems come in is we have systems that the design, either we've outgrown the design, or it was like quickly slapped together. Some engineer was like, I need to solve a problem this weekend for myself. They solve the problem, and everybody loves that solution. I actually did this for Akamai when I was there. I built a vacation tracking system for myself, and it became within two weeks the company's vacation tracker and stayed that way for 17 years.

A weekend project-- now, I maintained it. I used it for a lot of other things as well, but that's an example of a system-- it was not the most usable, but it was better than the paper-based system we'd had before it. And so we moved on, and now we're trapped in that world of like, this is the system. And so when companies are looking at they're like, what are our biggest challenges? That's where it always start. It's like, what's the system you just implemented?

And my favorite one is to pick on is Active Directory. No offense to Microsoft, who's not on this call, which is Active Directory is so amazingly powerful, but you have to configure it correctly. And so many people just implement it, walk away, and then they discover, oh, we got breached by ransomware. And the story looked something like, one of our machines got compromised, an administrator credential was available, and lateral movement happened. And now our entire enterprise is gone.

So that's your vulnerability was implemented a system that needed a little bit of design. And there's a safe way to implement Active Directory, and maybe 10% of companies actually do that.

I think, Amy, Andy's take is really interesting because, I think, anecdotally, people hear all the time that people are the problem. That you followed some link that you shouldn't have. I used to work somewhere where they were emailing us all the time like, this phishing email is going around. This is happening.

And you always thought, well, gee, if I'm getting alerted all the time, it seems likely that a few people will click on this before you get the email because by that time a bunch of people had alerted IT and so on. So yeah, give me your sense of that human vulnerability versus the larger picture?

Yeah, no, I completely agree with what Andy is saying. We should all. Education is a really important component about making sure that we're all acting in a way that's in our best interest long term. That definitely don't click on links, right. Verify through a separate channel any unusual requests you get. All of that's true.

But fundamentally, I think it's our system design that is really the key. How do we use modern architectures to build a system where the components that require trust are minimized in number? Where it is easy to verify their configuration to Andy's point about Active Directory and any other computer system that has ever existed ever. That it is configured the way that it should be.

And I think the partnership between product and security teams from the high-level design onward is such a crucial thing there because this is a really complicated and specialized discipline. It's also an organizational system designed to expect a developer to be up on all of the latest cyber research, right?

It's important when you're building anything to make sure that you have all of the different viewpoints together at decision time in the building of those products. I look at a stick, and I see a stick. My kid looks at a stick and sees I don't even know what. Like something totally different.

A sword probably.

Right, developers and security professionals it's the same thing. We look at a piece of proposed design. We look at a piece of code, and we see it from different perspectives. And that's how we make the system not quite as risky for the humans using it.

So Amy, can you build on that and just talk a little bit about where security fits right now, in general? As you can tell across the landscape like in the C-suite to what degree is it integrated into the core of a company and to what degree is it like an add-on or seen as some sort of specialization that's not core.

So I mean, I think it varies, both in terms of size of organization and the perspectives of the particular C-suite leaders. But also to Andy's point earlier sort of where in the maturity journey of the company is? But if we could get past the early startup hurdle where there is a C-suite. It's solid. There are mature practices.

I think we're in a place right now where nonsecurity leaders were maybe earning trust as security professionals in those moments. That we also want to see products released to customers. I think one of the things that I love most about my job right now is that we ask, how? Right, how do we do this? Does something need to be invented? Fine, let's figure out how to invent it. Maybe someone else has already inventing it. Let's figure that out too.

And I think the more security leadership in those executive conversations is talking about how, the more receptive the business leaders are to saying like, oh, yeah, I also want to protect our customers. We all care about our customers. Let's figure out how to do this quickly together safely.

And Andy, what do you see?

So I love the way Amy's phrasing that. Because what I often see-- I think the CSO is actually this weird artificial role in a company that originally showed up because CIOs stopped being about governance and enablement, and then started being about cost reduction and how do we hit the 80/20 rule. And so security started to being the, well, we'll do the security governance for the systems nobody else seems to care about.

When e-commerce showed up, it was never built by the CIO. It was built by some random application engineer and a security person tried to defend it. And what I'm seeing now is the CSOs and security teams moving of back into that role of enablement to say, our job isn't to keep the company from deploying unsafe things. It's to help the company deploy safe things. And it's a really subtle shift, but it completely changes how you approach the business.

I wonder from both of you just the very sort of basic question of, what do you feel like the biggest cybersecurity threat that companies face is, whether they realize it or they don't? Andy, do you want to take a crack at that.

So I think the biggest threat-- and it's a long-term little one. It's going to sound very odd-- is actually compliance. And the reason for that is it's really easy for someone to take this massively long list that they get, whether it's from PCI or the NIST 853 or FedRAMP, and basically, say, how do I do the minimum work to just do this stuff, and then I'm secure?

And so it's not looking and saying, what actually is my architecture? What are my unacceptable losses? What hazards expose me to them, and what do I need to do? It's just assuming this checklist would apply to you. And then you end up with, OK, you have a flat Active Directory infrastructure, you're not patching your windows machines, you have machines out on the internet that are connected to your databases.

Those are all just emergent problems that come from not actually having an architecture-based security design that looks at what your real risks are and deals with them. So it's not a specific technology risk. It's more this massive risk factor that is, how do I do the least amount to check the box that says I'm secure?

And probably because it takes a lot of work and money and person hours to do all that. Go the extra step, right?

It does. And you end up arguing with your auditors because they want you to do a thing that makes absolutely no sense. I can remember massive argument I had with the FedRAMP Program Management Office when I was at Akamai because they wanted us to do humidity measuring in every data center we were in.

Because FedRAMP is based on the idea that you're in one data center, so having humidity controls really matters if a data center could like every machine fries out. We were in thousands. I don't care. To me, data centers were like power strips. If one goes out, we just move to somebody else's data center. So I spent energy as the chief security officer of a multi-billion dollar business fighting with an auditor about the silliness of humidity controls.

Like, multiply that out, and when people get done with all those fights, do they still have the energy to say, hey, what are the risks that are unique to my business, or that wouldn't normally come up just by doing this checklist-based approach?

Yeah, just to build on what Andy just said, I agree. I think the biggest threat is a checklist approach versus a risk approach. Because it means that you struggle to adapt when things change, when new regulation comes out, when the product direction goes in a different place and you have to adapt to that.

If you've started with your understanding of your core risks, of your core workflows, of your core needs, and then built that in from architecture forward, you're much more readily able to adapt to change. But I don't know across the whole industry whether or not we're where we need to be there.

I think one thing I didn't hear people might be surprised is almost every day you see a huge hack of some hospital, museum. I mean, I can just think of so many examples. I wonder, why do these things continue to happen? Are they are they ramping down? Are they ramping up? Give me a sense of-- I mean, you read this news too. What do you see happening there? Put that in context?

So I think a lot of it is legacy architectures. That people are still using technologies that they deployed 10, 15 years ago versus what they could deploy today. And obviously, the day you deploy something is the safest you'll ever be with it because it's going to just become more unsafe over time. I think there is a piece of it, which is people and organizations haven't yet really made the mind shift to say, you have to stop trusting your administrators.

And I don't mean that don't trust the human. I mean, don't create this one spot where if someone gets access to it, whether it's a hostile insider, or an outsider who compromises an insiders credential, where they can do everything to you. And that almost always shows up in every one of these breaches somewhere in the middle was an adversary compromise that administrator and did X. And now they own everything. They have ransomware. They stole credentials. Whatever it looks like.

So as people do architecture-based design, that's the one thing we need to look for more is how do we stop trusting administrators with the keys to the kingdom as much as possible. Obviously, you'll never get them to zero, but right now, we basically give them infinite control, and that bites us whenever there's a breach.

Infinite control without then follow on side channel checks too, right? How do you build in other checks and mitigations and balances to verify that the person taking this action-- that there is a human that you trust that's taking this action?

Yeah, I mean, the technology I would love to see is in the endpoint defense. There's clients on every endpoint, and they're all basically a remote root access into that endpoint. So you install CrowdStrike or whatever you're going to install, and, all of a sudden, if you have an administrator breached, they have remote access to your whole enterprise.

How can we get to a world that we can have security technologies on endpoints, but do not trust anybody off the endpoint. That we can still verify that that's a safe machine but that we don't have any control of it from a central location. That'll go a long way towards protecting enterprises.

Andy, I wonder, what is your feeling about-- obviously, we see a huge amount of hacking coming out of certain countries like Russia, North Korea, China. It depends on exactly what they're trying to go after. And a huge amount of really impressive expertise actually attached to that. And I just wonder, is this a growing threat in mind? How you think about this?

So I think that threats it's been there, it's growing, but so are other threats. So it's hard to say. Is this growing faster than the for-hire adversaries. And we've certainly seen evidence of cases where Russian entities that were probably private enterprises then go and train governments elsewhere. And so, all of a sudden, techniques are being replicated across the environment.

That's a fascinating thing to watch. Not really great for us but it's always interesting. But the way that I tend to look at it is when you're looking at strictly information-based attacks. So not planting an agent inside your company. That's something that nation states are more likely to be able to do that nobody else can do.

But if what China's going to do to your enterprise is likely very similar to what a ransomware operator could do to your enterprise. They might have different motivations, and they might be willing to do different things. But, at the end of the day, you have to approach both of those from the same mindset of, how do I keep an adversary from gaining control of my environment?

Right. Right. Amy, we talked earlier about AI helping engineers write better documentation, for example. I also wonder if you worry about AI helping attackers to write better phishing emails, or AI can be used coming and going.

Yes. Absolutely. Absolutely, right. I think at the heart of any adversary interaction, both sides are making a cost-benefit analysis even intuitively. And AI makes certain things much cheaper to do both on the defensive side and on the offensive side.

I think the thing that keeps me optimistic is that we do have these opportunities to use AI defensively to pattern match, to prioritize, to identify, and then focus the human's attention on the places where it's most likely to be needed.

I wonder from both of you when you look ahead, what is the-- I wanted both. Maybe Amy, we can start with you. Like, what's the one thing that concerns you the most about what's coming down the road, but maybe not here right this second, or maybe, at least, not most noncyber people aren't thinking about it, and then what's the thing that makes you most an optimistic and forward?

Sure. I think the thing that worries me most is our ability to adapt to the pace of change. I think we're at the start of some really great things in the security field, how we work as a core part of the business, how we're building products. But we've been developing products for a long time now. And they weren't necessarily all across the world developed with that mindset.

And so our ability to adapt is the thing that I'm most focused on changing. And I'm optimistic there because I see so much more partnership opportunity. When we sit down as security partners in a high-level design meeting, when we look at trying to figure out the tooling required to make the fastest code, to write the most secure code to write, I see a lot of opportunity for technology and research to support that transition.

And Andy, what do you think?

So the thing that worries me most is actually still going to be email. It's sort of the last federated protocol we have where anybody can throw up a client or a server. We run basically the world on email and is the most unsafe protocol out there. You keep coming back to phishing. Very real threat. But why does phishing work?

Because I could send email and claim to be you, and your mail browser will tell you you just sent mail to yourself. Like, that's a problem. And then you click a link, and bad things happen. And what we're seeing is everybody's bringing out their own internally unified communication system, whether it's Slack or Zoom or Teams or Webex, they have their own. They're not federated as well, so they're not interacting.

And for me, how do we run businesses that currently have a command and control infrastructure-- there is email-- in a way that's actually trustworthy and safe. That when you get a message from your CEO, you know whether or not they wrote it and not that the human has to do forensics to figure out whether or not this actually came from the CEO.

And is the system that you dream of, is that coming?

I would love to say that it's coming, but I don't see that happening right now.

OK. I just want to get Amy to chime in. Do you worry about email too?

Yeah. Sure, everything that Andy said is true. I'd like that system too. There are lots of ways, though, that you can build in-- back to the context and the holistic way to build products-- there are lots of ways that you can adapt your workflows, such that particularly risky or high impact communications get verified by an outside channel. Would it be great if the email protocol and systems built on it were able to make those unnecessary?

Sure, but it's not-- if you get a really unusual request, it's not like we don't also have a phone to call and say, hey--

I was going to say--

--a really weird thing, right?

Yeah.

There's no way to double check it, either you show up in person, or yeah, you call them on the phone. There's these double checking things.

Right. If you're moving from a checklist approach to an understanding what the risk is, and then making sure that your whole system addresses that risk. You're in a much stronger place. And I do think we will given the impact of AI on great phishing email, I do think we'll see a normalization of checking in other channels to make sure that something that seems a little strange is actually legit.

Yeah, the best thing that CEOs can do is actually do that normalization. Many of the CEOs I work with have started a thing where like-- because there you get the gift card, the text messages-- you get a text message claiming to be from the CEO, I'm in a meeting. I need you to get gift cards. And the CEO has normalized take that, screenshot it, and drop it on Slack, and tag them and the CEO will say thank you.

And they'll comment about, oh, yeah, it's a bad quarter. We need you all to buy gift cards. Ha-ha-ha. But you have to recognize that every attack is an awareness opportunity. That if 90% of your employees are doing the right thing, they need to do it more publicly so that 100% know what the right thing is.

Andy Ellis is operating partner at YL Ventures. Amy Herzog is chief information security officer for Ads & Devices at Amazon. Thanks so much to both of you. This was great.

Thanks.

Thanks.

[MUSIC PLAYING]

If you want to know more about the CSAIL Alliance's program, and CSAIL latest research, just visit our website cap.csail.mit.edu. I'm Kara Miller. Our show is produced by Matt Purdy and Nate Caldwell with help from Audrey Woods. Tune in next month for a brand new edition of the CSAIL Alliance's podcast and stay ahead of the curve.