MIT CSAIL Alliances | Vinod_Srini_Project_6

Welcome to MIT's *ComputerScienceandArtificialIntelligenceLabsAlliancesPodcast.* I'm Kara Miller.

[MUSIC PLAYING]

Today, we build on last month's discussion with cybersecurity leaders in business and take a look at the future of online security.

Creating a central point of failure is or will become, at some point, a problem. It's not an if, but when. And when that happens, you know, I don't know, billions of dollars worth of data will get exfiltrated. And then there's nothing we can do about it, really. Plus, quantum computing is on the horizon. And that could totally and completely change the game.

Our most sensitive data, you know, someone could be actually vacuuming it up at this point, hoping that in N years, they will have a computer that's powerful enough to decrypt it.

But even today, without quantum, folks in the trenches have plenty to worry about, including email.

If you are in a situation where you say, I'm only going to get email from people I trust, well, you're going to miss a lot of emails.

That's right.

That's a fact.

That's all coming right up.

[MUSIC PLAYING]

On our most recent show, I talked with two executives in cyber security and asked, so what are your biggest problems on the ground right this second? And since we've got two folks this month who are looking into and creating the future of cyber, let's start off by tossing them some of those executive concerns and seeing what they say.

The first concern comes from Amy Herzog. She's Chief Information Security Officer for Ads and Devices at Amazon.

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-- like, 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.

So can we adapt better? And how would we do that? Here's Srini Devadas, a professor of electrical engineering and computer science at MIT and a cybersecurity expert.

What's happened is the scale of the number of people on devices that are connected is clearly, you know, increased dramatically over the last five years and continues to increase. And I think that adaptation comes from two things. One is what I just mentioned. But the other part is that the services that people want to use, you know, the variety of those services has increased as well.

So I go back to-- yes. I mean, you know, absolutely. I mean, it's a challenge. But technologies and design principles, you know, do exist to solve these hard problems. And I think there's an educational component here with respect to the CISOs' understanding of, you know, all the technology that's out there that maybe hasn't quite been deployed.

But, you know, don't wait for the Amazons or the Googles to deploy it. You can do things at a smaller scale if you're a smaller company. And I do think that there's a ton of great work that could solve these problems, but it stays in the research oven.

Vinod Vaikuntanathan, a cryptographer and professor of computer science at MIT, says yes, adaptation is really tough. And you need a multi-pronged approach.

We need both short-term measures and long-term measures. We need to think deeply about how these tools affect security as it is practiced today and, again, once again innovate. You know, and I agree with what Srini said, which is there are a lot of ideas from the academic realm, some of which, you know, probably are not interesting enough yet to be in the real world. But some of them are.

And I think it is important to have a sort of a two-way dialogue between the CISOs and academics, both for us to understand how the problems that they are facing are changing, and for them to understand, you know, what the tools are that we are thinking about and developing.

OK. Now, let's turn to a different concern. Andy Ellis, the former chief security officer at Akamai, worries about bad computer code popping up everywhere. How do you deal with it?

How do we train AI to write safe code rather than write more of the code that's already been written? One of the things almost every security person does now, or has done at some point, is you count how many pages into any textbook on programming is the first vulnerability. And it's usually, like, the first time they're introduced, you know, in taking inputs. It's like we teach people to write unsafe code. Can we fix that?

Devadas says he doesn't see why not.

You want to teach AI models to learn from other people's mistakes to avoid them, right? I mean, which is really avoiding bugs.

And I heard about this wonderful class at Arizona State University that's a graduate class in computer security, where the professor was using LLMs to help his students attack vulnerabilities, attack pieces of software that had known bugs, and go through the process of discovering the vulnerability.

And so-- I mean, obviously, you know, knowing the vulnerability, discovering the vulnerability, is step one to writing safe code. So I think, uh, absolutely. I mean, I think the training data needs to be collected.

We all know that these deep neural nets need lots and lots of training data. But I think once you have the training data, you should be able to defend against a large class of vulnerabilities.

Now, they have these things called zero-day vulnerabilities which no one knows about. Clearly, they aren't in the training data. It'd be super cool-- I don't know how to do this, I'm not an AI expert-- to write safe code that is protected from zero-day vulnerabilities.

So, you know, just to ask a different question, perhaps easier than what Andy was asking, which is can the machines actually tell if a given code has vulnerabilities or not? You know, can we use machines to detect vulnerabilities, potentially zero-day vulnerabilities that we don't have-- that are not in our database quite yet?

Or potentially, like, that you're using AI to get the code, but also AI to sniff out where the problems are in the code.

I suppose it's a complementary question, but yeah.

OK. But it sounds like it's also a question that has not been settled or solved.

To the best of my understanding, yes.

And one final concern from our cyber executives, and this is a big one, here is Andy Ellis.

So the thing that worries me most is actually still going to be email. It's the last federated protocol we have, where anybody can throw up a client or a server. We run basically the world on email. And it is the most unsafe protocol out there.

Like, you keep coming back to phishing-- very real threat. But why does phishing work. Because I could send email and claim to be, 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, you know, everybody's bringing out their own, you know, internally unified communication system, whether it's Slack, or Zoom, or Teams, or Webex. Like, they have their own. They're not federated as well, so they're not interacting. And for me, like, how do we run businesses that currently have a command and control infrastructure that is email in a way that's actually trustworthy and safe?

So, email. How insecure is it? And can it get more secure? Devadas and Vaikuntanathan say it's not that secure. And getting it to be more secure won't be easy.

You know, authentication of email is a big problem. People have been trying to solve that. It's such an open system. If you are in a situation where you say, I'm only going to get email from people I trust, well, you're going to miss a lot of emails.

That's right, that's right.

That's a fact. So how do you set things up so the first time you get email, you know, it's created in an untrustworthy way, or are treated-- that potentially untrustworthy, and then after that, maybe you grow the trust, right? And there's definitely proposals that have been put out that do things like this, and some of them have been deployed.

But interoperability is a huge issue. You know, you just want to be connected. And there are people who lose business if they don't respond to emails, if they don't click on links, right?

And then there's the people who don't know any better, right? Because it's kind of obvious, but they end up clicking anyway. Email is the oldest internet application. And it's also the hardest to fix for, I think, kind of the same reason, right?

But I'd say that the closed systems, if you think about enterprises and you think about, you know, when you get messages from people, you see external email, untrusted sender, you know, sometimes when they respond to you, well, that's the closed system at, you know, Amazon or Singapore defense organization or what have you, tracking, you know, all the incoming email. And I guarantee you that they have fewer problems with phishing attacks.

And so, you know, one way I think is levels of trust. Things that come from your enterprise-- sure, click away, right? But flag with colors or what have you-- and also virtualization.

You know, run-- when you click on something, you're not clicking in an environment where you're giving access to private data, but you're clicking in kind of a test environment, right? You know, it's a little bit of taking the bomb into a, you know, into a secure--

[INTERPOSING VOICES]

And it's kind of scary to say that, you know, for every email that you get. But, you know, as long as it's seamless and you can hide that from the user, I think it's doable.

So just separating the problem into two parts, right? I mean, protecting emails for the general population seems like so much of a harder problem than protecting emails within an organization. Because you can build in, you know, in principle, all these sandboxing strategies that, again, should work. And one should be able to do it.

For the harder problem, you know, what Srini said is really interesting, you know? If you think about the way we interact with people in the real world, you know, I meet a stranger, you know, I'm a little bit skeptical, right? And we build trust over the course of many interactions.

And that is not something that happens with email, right? I mean, so I had I have a colleague who once tried to implement this strategy, which is anybody who sends him an email has to solve a math problem.

Oh, my god. [CHUCKLES]

So that's one way to filter it out.

So we only had mathematically inclined friends, that's fine.

But-- but, you know, that's-- or solve a CAPTCHA or whatever, right?

Right, right.

So probably not perfect, but sort of, you know, an add-on-- Srini's add-on would be the first time you talk to me, you have to solve a mathematical problem, but maybe later on. It's fine. So I think sort of incrementally building trust is a notion that I don't think exists as much as it should, you know, in these email systems.

So what problems do academics who work on cyber see coming down the tracks? What worries them? Well, one of the big problems is with reporting on what goes wrong, which they say is something we don't do nearly enough.

Interestingly, there are industries that have used robust reporting to enhance performance. One of them is aviation. Back in the mid 1960s, there was a man named Bobby Allen, who at the time ran the Bureau of Safety at the US Civil Aeronautics Board. And he had what turned out to be a pretty brilliant idea.

He argued that people didn't really want their names to be attached to on-the-job mistakes. But if there was some way of looking at all the data on those mistakes without getting anybody in trouble, well, that would be progress. He was talking about coming up with a better reporting system for aviation, a system which now exists. It has been held up as an example for industries like health care.

But you could argue aviation has a lot to teach cybersecurity. Too often, say David Austin, Vaikuntanathan, openness is not the order of the day in cyber. And just as prolific accident and near-miss reporting in aviation helps everybody in the field, Vaikuntanathan argues that the lack of that sort of information in cybersecurity, well, it can be a problem.

A consensus that I hear from a lot of the CISOs and people in power is that we need to be more data-driven in understanding the attacks that happen, right? So at this point, you know, a lot of these cyber attacks, they happen. You know, people internally within an institution have data about it. But there is no data sharing of any form.

So I think people make mistakes. But then if you make mistakes, you have to learn from them, right? Otherwise, you know, we are doomed to fail again and again and again.

But he says there are some green shoots, important work that's going on in this area.

We have a project here at CSAIL called SCRAM. This is with the Internet Policy Research Initiative and folks from Sloan. And the point there is to actually build a cryptographic system that allows us to collect these kind of data from different organizations while preserving the privacy of this data.

So these organizations, you know, they don't need to be worried that they're actually sending us the data. They don't need to trust us. They can encrypt this data. And what we can do is we can run these statistical computations that figure out various facts about cyber intrusions on the encrypted data while it stays encrypted. So we don't see anything, but we can still use the collective intelligence and the collective data.

Of course, back in the 1960s, commercial aviation was still a fairly young industry. And the realization that data could help it solve its problems, that was key. As more of our lives migrate online of course, data has become central to protecting ourselves. And in 2023, according to the FBI, losses to Americans due to cybercrime topped $12 billion.

You could argue that one positive development over the last few years is the centralization of privacy efforts. Lots more CEOs now have the option to use Amazon Web Services or Google cloud, for example. So is that a good thing? Well, it depends on who you ask. Vaikuntanathan isn't so sure.

So I'm a fan of distributing trust. You know, of course that means that you have to enable multiple organizations or centers to be competent enough to manage the data. So that definitely is a trade-off. But creating a central point of failure is, or will become at some point, a problem. It's not an if, but when this will happen.

And when that happens, you know, I don't know, billions of dollars worth of data will get exfiltrated. And then there's nothing we can do about it, really. So I think it is a good idea to try to think about implementing ways of distributing trust.

There are ways from cryptography, from the principles and foundations of cryptography that enable you to distribute trust. The question is to what extent they will end up being used in the real world.

So just for example, you know, there is a company that actually takes digital signature keys from Bitcoin wallets and sort of shards it, in a sense, so breaks it up into many pieces, and distributes it to many locations. So if somebody gets like two pieces out of five, they won't be able to really use it. So that sort of technology is there. The question is to what extent we will see it in the world.

Yeah, so I'll show my bias here with respect to the type of technology that I would like to see more widely deployed, which is secure hardware, where you need a physical attack as opposed to a remote software attack to discover secrets. And there is technology out there. Intel has processors that have these extensions called software guard extensions-- and originally deployed in the Skylake processor and subsequent generations-- where you essentially tunnel in, you know, into a piece of hardware, you know, an integrated circuit chip that has a secret key embedded in it. I mean, it's a public-private key pair. The private key needs to be protected. The public key obviously can be public.

And think of it as SSL into this machine. You know, and then the only place that your data is decrypted is, you know, voltages and currents inside of this chip, right? And any time you leave the chip, you know, things-- you use cryptography. You know, you do use cryptography, things that are encrypted when you are putting things in memory. I mean, not just disk.

I mean, people realized, you know, maybe 30 years ago that, hey, you better encrypt things on disk, right? Because people can walk away with a solid state disk or a magnetic disk. You can do cold boot attacks and people can walk away with memory, you know, Digital Random Access Memory, DRAM.

And so now these machines like Skylake, they encrypt using these hardware keys, encrypt data that's in memory, you know, your gigabytes of memory. So I think it is a central point of failure. I will say that. And I think that's-- you know, it is a bit of an issue. I would agree with Vinod. But it's a central point of physical failure. That's the way I think about it.

And, you know, you can have bodyguards by the gazoo in front of the data centers and, you know, check, you know, biometrics and cardkey access and things like that. But this is not the only way of doing things. I mean, you could do this and you could do what Vinod said about distributing trust. And then maybe, you know, we got security heaven, right?

Vinod, let me dip into quantum. And before I ask, let me just back up for a minute. I've talked to people who feel like quantum computing is imminent, kind of around the corner. Other people who feel like, no, no. This is, like, barely even worth talking about right now. So first, just give me your take on that, because I know you've been thinking about cybersecurity in relationship to quantum computing.

Right, so I'm teaching a class on quantum cryptography. So we try to come up with different possible worlds that people could live in. There are people who live in the no quantum world, as in, we are all-- I don't know. What is there to see here, right? We're all classical. Versus the world of quantum [INAUDIBLE], where, you know, our iPhones will, in 20 years, have quantum states in them.

So these are two polar extremes of how people-- certain people think. So at this point, you know, I think it's unquestionable that there are real advances being made in building quantum computers.

Now, you know, in terms of cryptography, I think the question we should be asking is, what is the chance that in 10 years, we will see a nation state with reasonable quantum computing power to break factoring and RSA-based systems and discrete log-based systems? So the point being that we can't sort of, you know, be complacent and expect to wait until that point of time and then scramble to change our, I suppose, internet infrastructure. That's just not going to work.

The calculus, though, that I think people are doing, and I've spoken to several people from the industry about it, is they're trying to estimate, what is the expected amount of time until when we'll have actually a scalable quantum computer? Is it really infinity, you know, never going to happen? Or is it, like, five years, right?

So it's really very hard to say because I'm not involved in actually building a quantum computer. But I would say that we have to anticipate this event, this sort of singularity, by, you know, at least five, maybe ten, maybe more years.

And once it happens, there's potential serious problems for the way we currently encrypt things, right?

That's right, that's right. That's exactly right. And the kind of thing is, you know, we may not actually even realize that this is happening because, you know, our most sensitive data could be-- you know, someone could be actually vacuuming it up at this point, hoping that in N years, they will have a computer that's powerful enough to decrypt it.

The calculus is complicated. Because you could ask, what kind of data do I want to protect for the next five years, you know? Probably not my text messages, you know, to my family. You know, that doesn't seem like-- maybe my financial information? Unclear. Maybe government secrets? That seems to have a longer lifetime.

So it's a bit of a complicated calculus. And I think people are still trying to figure out if they should move to systems which are presumably quantum secure. And in fact, NEST has recently standardized a few such systems. And people are kind of actively thinking about whether they should-- you know, how fast they should adopt it. Yeah, it's a complex question.

You talked about-- you mentioned nation states. And I wonder, when you think ahead and what the future of cybersecurity looks like, whether you feel like the biggest threats on the horizon are, in fact, nation states, who obviously have a lot of money?

And also, not just money, but like long-term thinking at their disposal and the ability to invest in things over time and hire people and keep them employed, let's say, for the long run, and some very talented people. Or that's not really-- you know, it's more like individual, rogue, non-nation state actors that we should be thinking about. I wonder what your view is of 10 or 20 years from now.

What scares me is a nation state with a whole bunch of really smart people who are going after critical infrastructure that, you know, obviously has a cyber component. That's what scares me. So it's a little bit of a combination of what you said, Kara.

So like they turn off the lights in San Francisco for a long time or something.

And simultaneously, they're doing other things in Manhattan.

For sure.

You know, it's the-- they're also distributed across the world. It's a distributed denial of service attack. It's a distributed intrusion. It's a powerful attack. It has variety. It has subtlety.

They are holding back on zero-days that aren't really zero-days from their standpoint because they knew about it months ago. But they are accumulating them. And then it all comes to a head. So that is a movie. You know, that's a scary movie.

It is, for sure. It's a weighty but important note to end on. Srini Devadas and Vinod Vaikuntanathan are both professors in the Department of Electrical Engineering and Computer Science at MIT. Both are part of CSAIL. Thanks so much for being here.

Thank you.

Our pleasure.

[MUSIC PLAYING]

And if you want to know more about CSAIL's latest research and the CSAIL Alliances program, please 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 Alliances Podcast* and stay ahead of the curve.

[MUSIC PLAYING]