ICANN86 Seville | PF – Current and Coming Impact of Artificial Intelligence on DNS Abuse
Monday, June 08, 2026 – 16:30 to 18:00 CEST

KATHY SCHNITT

Hello and welcome to this session on Current and Coming Impact of Artificial Intelligence on DNS Abuse. My name is Kathy, and I am the participation manager for this session. Please note this session is being recorded and is governed by the ICANN Community Code of Conduct, the ICANN Expected Standards of Behavior, and the ICANN Community Anti-Harassment Policy.

Interpretation for this session will include English, French, and Spanish. Please observe the following guidelines to participate in this session. I will also post this in the chat for your reference. Only questions posted in the Q&A pod will be read aloud during the session as time permits and when directed by the chairs of the session. If you wish to speak during this session, please raise your hand in Zoom. Use your table microphone to speak. If you do not have a table microphone, when we call your name, raise your hand, and we will come to you with a roaming mic. When speaking, please state your name and affiliation for the record, the language you wish to speak if speaking any language other than English, and maintain a moderate pace.

With that, I'm happy to hand the floor over to BC Chair, Mason Cole, and SSAC Chair, Ram Mohan.

RAM MOHAN

Thank you so much. For years, our discussions around DNS abuse have centered on a familiar list. Phishing, malware, botnets, and other forms of malicious activity that exploit the DNS. Those threats are still with us, but something important has changed. We are witnessing a fundamental shift, not only in how abuse is created, deployed, and scaled, but also how it is detected and mitigated.

Artificial intelligence is changing the game. What once required deep technical expertise can increasingly be accomplished with AI-assisted tools. Attackers can generate convincing content in seconds. They can build infrastructure at scale. They can adapt their tactics in real time. The result is something new. Machine speed abuse. Campaigns can now be launched, modified, and scaled faster than human teams can reasonably respond. For the first time, we are confronting a future in which attackers may be able to innovate and operate at a pace that consistently outstrips traditional defensive approaches.

Defensive users of AI are also still in their early stages. The race between attack and defense is only beginning. The question is no longer whether AI will affect DNS abuse. The question is how quickly and how defenders can keep pace. So the stakes are high. This is not simply a technical challenge for registries and registrars. It is a challenge for the entire internet ecosystem, the ICANN ecosystem that operates in there, from technical operators to

security practitioners to policymakers, governments, businesses, and ultimately the users who depend on the DNS every day.

Today's session is designed to help us better understand and prepare for that reality. To understand what this means for the DNS ecosystem, we will focus on three areas where AI is already beginning to reshape the threat landscape. First, the democratization of attack and mitigation tools: how AI is placing increasingly powerful capabilities into the hands of actors with limited technical expertise, and how defenders are responding. Second, the rise of machine speed abuse: the automation and scaling of malicious activity at an unprecedented speed and volume. And third, the role of coordination: what ICANN and the broader community can do to strengthen awareness, improve information sharing, and build more resilient approaches to mitigation and defense.

This is intended to be an interactive session. We have experts joining us from technical, operational, and policy disciplines, and we encourage your participation throughout. As Kathy said, if you have questions, you will need to be logged into Zoom to raise your hand and ask questions, but we will also be using audience polling and open discussion to bring community perspectives into our conversation.

Our goal today is not to predict the future with certainty. It is instead to understand where the threat landscape is heading, and to begin thinking together about how the ICANN community

should prepare for what comes next. Thank you for joining us. Let us begin. First, to the Mentimeter slide that you have up there, please use the QR code or the number code to join. While you're doing that, Mason and I are joined by four panelists, and I apologize in advance if I don't pronounce all names perfectly.

To my left, I have Gabriel Andrews. Gabriel is a 15-year veteran of US law enforcement with a focus on criminal computer intrusions and cyber-enabled crime. He is currently an SSAC member and co-chair of the Governmental Advisory Committee Public Safety Working Group.

Rowena Shoo, on my right. Rowena is the Senior Director of Industry Affairs and Policy for the NetBeacon Institute. She previously worked for Ofcom, the UK Communications Regulator, Nominet UK, and the UK Government at the Department of Digital, Culture, Media & Sport. Rowena has over a decade of influence in and around policy and holds two degrees: Bachelor of Laws and Bachelor of Arts from the Australian National University. And Rowena, I didn't mention the honors or the international relations and political science as I was doing it.

We also have Samaneh Tajalizadehkhoob. I don't know how well I got that. Samaneh is Director of the Security, Stability, and Resilience Research Team in the office of the CTO, OCTO, at ICANN. She holds a PhD at the intersection of web security and machine learning and acts as OCTO's liaison to the SSAC. With more than 12 years of experience, Samaneh's work focuses on designing data-

driven metrics and research frameworks that translate complex security challenges into actionable and measurable metrics. Samaneh is behind multiple ICANN projects such as INFERMAL, DAAR, Metrica, [SEDAG], and the Tranco popularity list.

I also have Laurin Weissinger. Laurin is a member of the SSAC and one of the chairs of the SSAC's AI Working Group. Outside of ICANN, Laurin consults the National Institute of Standards and Technology (NIST) on AI evaluation and teaches cybersecurity at the University of California, Berkeley, and Tufts University. Previously, Laurin worked hands-on and in more management-focused roles in industry and academia. He holds a DPhil focused on security and digital trust, and contributes to various anti-abuse working groups.

So, with this introduction, let me hand this off to my co-moderator, Mason Cole.

MASON COLE

Thank you, Ram. Good afternoon, everyone. My name is Mason Cole. I'm chair of the business constituency, and it's a pleasure to welcome you all here this afternoon. Let me add my welcome to everyone. My thanks to Ram and to the SSAC for co-sponsoring this session. The BC, as many of you know, has long been interested in and has advocated for a robust ICANN response to DNS abuse, as it significantly impacts the businesses we represent and their users and customers. We're grateful to the SSAC and to the community for sharing our interests.

The focus for today is that this session is meant to be informative and not prescriptive, meaning we're here to introduce a topic of reintroducing a developing issue of importance to the community in the way of observations and information. Neither the BC nor the SSAC, as co-sponsors of this session, is proposing policy at this stage. Once our panelists have concluded their presentations, we hope for a robust discussion here in the room about AI and how it may impact the DNS, and we hope this sets the stage for further cooperative discussions in the community as we continue our mutual work on DNS abuse.

So please, again, use the QR code or the number code to join. While you're doing that, I'll set the stage for our panelists, and that is to say that each panelist will deliver a high-impact, time-limited opening statement anchored to the three key focus areas, which are: the democratization of attack tools, the move toward machine speed abuse, and the potential for ICANN-led coordination in developing AI-resilient defenses. Laurin will spend his four minutes focusing on academic research findings, how the advent of AI shifted the fundamentals of both abusers and defenders, and how the threat is expected to evolve over the next three to five years.

Samaneh will focus on SSR research, what the impact is of AI on the actual observed cases of DNS abuse, and what ICANN's role is in monitoring and responding to the rapid changes in DNS abuse amplified by AI. Rowena will focus on the industry's realities. That is what's happening in the real world with operators, where the challenges are, and how AI is currently being used in DNS abuse

campaigns. And finally, Gabriel will focus on the impact on law enforcement, making the threat tangible and explaining how AI has become a game-changer. What does machine speed abuse look like in real life? And he'll provide actual examples.

Let me hand it back to Ram for the introduction of Theme 1. Ram?

RAM MOHAN

Thank you so much. Let's begin with the fundamentals. When we talk about democratization of attack tools, what do we actually mean? For years, sophisticated DNS abuse required specialized skills, significant resources, and substantial infrastructure. Today, AI is changing that equation. Capabilities once reserved for well-resourced actors are now becoming accessible to a much broader population. The barrier to entry is not just falling; arguably, it is disappearing. But this is only half the story. The same technologies that empower attackers can also strengthen our defenses. AI has the potential to help the defenders identify threats more quickly, analyze patterns at scale, and respond more effectively.

That brings us to a critical tension. How is AI being used in DNS abuse today? How is it being used on the defensive side? And as these capabilities become increasingly accessible, will defenders be able to keep pace? Or will the economics of this emerging arms race increasingly favor the attacker? To help us understand the current reality and where it may be heading over the next several years, let me turn to our panel.

But before we get there, we should also get to the opening statements for each of the panelists. So let me go to the opening statements first, and then I'll come to each of you on the democratization questions. Thank you, Mason, for that. Laurin, let me start with you and your opening statement.

LAURIN WEISSINGER

Absolutely. Thank you, Ram. Hello, everyone. I am looking at the Mentimeter slide here, and I'm seeing that people are concerned about increased speed and scale, and more convincing phishing and impersonation attacks. That works really, really well because my example was to talk a little bit about spam and phishing detection as an example, because it really represents quite well, I believe, the overall state of the research and academic literature. And that looks as follows.

You'll find one paper that says, "Oh, we used Claude," or "We used OpenAI tools." Just the LLM as it is, add a chatbot, and let's see how well it can classify phishing or spam emails. And, you know, there might be a relatively good output coming from there.

Then you move on to the next paper, which then says, "Well, you know, we did not get the accuracy that these other authors did, and we also really question how this would be usable or doable at scale." This really speaks quite a bit to the research landscape overall, where you have signals that all speak to tiny pieces of the overall picture of AI, where certain things work, and certain things

don't. We cannot always fully figure out, okay, why is it working here and why is it not working there?

That said, obviously, the use of artificial intelligence is nothing new in the DNS abuse world. But going from relatively simple automation to using machine learning -- and it has been going on for a long, long time -- moving now into, okay, how can we leverage, for example, LLMs for this last step of dealing with abuse?

So, where are we then with this? And I think this is really a little bit of a, how can I put it, like a little bit of a break on everything. We still have a lot of things we don't know when it comes to AI more generally. Also, therefore, for DNS abuse and what's going on within the remit of ICANN and the DNS ecosystem. Nevertheless, there is no question that either we are already seeing that the speed and scale are increasing, or we can clearly see that if we don't have the signals yet, the capability to do so is likely there.

As Ram was describing, you don't have to go to the highest level there is using AI. But the uplift, as we often call it, does exist for moderately-skilled or even lowly-skilled attackers to really increase what they can do. Easy example, stuff like translation, stuff like using LLMs to make stuff look more legitimate. And that raises the water level on the defensive side, trying to detect those things. Like in the past, things like, oh, this looks exactly like all these other emails. That was a signal. That might be going away or is more easily to get rid of.

Now, at the same time, there are really powerful uses of artificial intelligence on the defensive side. We can particularly employ artificial intelligence when we don't have these clear signals. And that is very common in DNS abuse. Whereas, not a "Oh, if this flag is set, it's definitely bad. If it's not set, it's definitely good," when we don't have these binary decision criteria in a lot of cases. So here it is really useful to be able to bring in multiple data sources and use heuristics more than very specified formal logic.

Allow me one more zooming out before I hand over to Samaneh, and that is that when we zoom out a little bit more, and we talk about our understanding of AI as a whole, not for DNS, not for DNS abuse, the same thing applies. There are a lot of things we do not understand as an AI community in terms of how things work, why certain things happen, or don't happen. And that includes, as far as I'm aware, as far as they communicate, the frontier AI labs. Things like, well, why does it help us to push some randomness into this model? And what does it mean for us to apply something seeded in randomness into queries for how predictable or how usable a system is? It's just one of many examples where it goes further.

As Ram was saying, I work mainly on testing and evaluation of AI systems. And here you have the same question, where it's like, well, if we don't really know how to figure out what these things do and how to evaluate them well, there are a lot of questions regarding their use and their impact going forward.

RAM MOHAN Thank you, Samaneh.

SAMANEH TAJALIZADEHKHOOB Thank you, Ram. Thank you, Laurin. Today, I'm going to shed some light on the perspective of AI and DNS abuse, in particular, the research perspective. I would like to have some definitions here of what we are talking about when we say AI, because to me, it sounds like all different things. I'm highlighting two particular things. One, the first, the large language models that help in creating content or facilitating creating content and bypassing language barriers. The second is AI agents that are particularly autonomously executing tasks that can register domains, create infrastructure campaigns, and do things on their own.

This is what we see in terms of how DNS abuse changed. In short, if you look at the blue area, we do not see any particular increase in terms of the total counts that we can say AI is changing the landscape or is causing a lot of increase in what we count as DNS abuse. However, we do see, if you look at the red line, that certain trends are changing.

For example, the red line here is showing subdomains -- accounts on subdomains -- which fall under the public suffix list. We speculate that this is related to domain names. When we look at the subdomains, this is related to domain names that are particularly recommended by large language models for different

development works. But in short, we do not yet particularly see any massive changes in the trends over the years.

Given that, AI as it is now is not fundamentally changing what DNS abuse is. What we foresee changing is the economics of abuse. So basically, lower barriers to entry, less technical expertise needed, language skills, and operational resources. Also, how the internet infrastructure is created, used, and abundant.

If we take that as an observation for the current state of DNS abuse, then our focus as a research team is on making sure that we are correctly detecting and measuring abuse. So, basically, not what abuse is, but making sure that our measurement methods are as dynamic as the space is evolving to be able to measure the correct stuff.

We, as a research team within ICANN, want to remain conservative on what we suggest. We want to remain suggesting things based on facts and not anecdotes. What we are mostly working on at the moment is developing measurement frameworks, data sets, and indicators that adapt to what we see. For example, we do see some changes in domain generation algorithms and in the subdomain space, and also the uptime of domain names. We adopt our measurement methodologies to detect those, but also remain conservative enough not to state things that we do not see in our data yet. That would be it for me. Thank you.

RAM MOHAN Thank you. Over to you, Rowena.

ROWENA SHOO Thank you very much. I probably didn't need my title slide, as you've introduced us all. I will just say the NetBeacon Institute is a part of Public Interest Registry, in case anyone wasn't aware. I think, as we've been talking about already today, no one really knows the exact impact yet of AI on DNS abuse, but it does seem very likely that it has the capacity to increase the acceleration of the scale and the quality of the attacks.

For example, writing phishing emails, vibe coding websites, creating convincing images, or using personal data that's available to create a really targeted attack. But what's interesting is when we look at our data that we have at the NetBeacon Institute, we don't see this huge acceleration or explosion that you might expect, given the capacity that we've been talking about.

You can see in this chart this is unique domains used for phishing over the time period that we've been collecting data. And you can see that there's a big peak in March 2025. That was our highest month on record. Most of the domains we saw in that peak were maliciously registered, but they were highly concentrated. Over 50% were concentrated in one registrar credential. I think what we do find when we dig into those increases and those peaks is we do find automation, and we find campaigns. But we do need to be aware that automation is possible without the use of AI, as well as with the use of AI. It can become quite difficult to really tell how

much of an impact AI is having, yet from the perspective that we're sitting.

This is just an example of what I mean by automation and campaigns. Often in the data, we see campaigns with a particular target, and there'll be a large number of unique domain names that use similar patterns and look to be algorithmically generated. This one in particular was targeting the UK government. You can see some examples of the kinds of domains that I'm talking about, the use of keywords, and the patterns that appear here.

This looks like automation, but it's difficult to tell whether it's the case of AI or if it's a case of the on-sale of customizable kits for phishing. So, phishing as a service, which we know from security researchers, happens. I've just gone over the first point, which is there is this capacity for acceleration, but it's really tricky to understand if that's happening in the data that we have. I think registrars and registries are in a similar position where they see campaigns, but it's difficult to know whether AI is driving those.

The second point I wanted to make is around AI also helping defenders. I think there is a huge capacity here for better evidence collection, for language and context interpretation, for vulnerability detection if you're thinking outside of just malicious registrations.

But I do think we need to be careful and thoughtful about how this is implemented because AI should usually be combined with some level of oversight from humans. The accuracy and the quality of the

decision-making are important to assess for the task you're putting it to. We also want to consider any impacts or risks if we're involving any external systems in our processes. Hallucinations and overconfidence are real risks if we're trying to use AI as part of DNS mitigation.

Finally, what's the impact on policy? Often, people talk about how technology moves so quickly, and policy needs to keep up. This is always going to be a factor with how we operate, but that doesn't mean that we need to give up on making policies to solve problems. We should remain focused on the principles of those policies.

For example, in PDP 1, on abuse that's taking place at the moment, the real heart of this principle is around unearthing large-scale malicious campaigns based on one well-evidenced report. In the second PDP that's been planned, it's really about applying friction at the point of registration to reduce the likelihood of taking in malicious registrations. We can't expect one policy to solve all of our problems, but we can, I think, by addressing targeted, specific issues, we can make real progress by staying focused and moving forward in small, solvable spaces. Thank you.

RAM MOHAN Thank you so much. Gabriel.

GABRIEL ANDREWS

Hi, so my name is Gabriel, and I felt that my role here would be most useful if I were able to bring a specific example of where some of my colleagues in law enforcement have seen AI employed by criminal actors and how they're using it. I will note that this example deals with ransomware -- initial access brokers. Those are the guys who first compromise a victim so that that victim can then be infected with ransomware.

Sometimes the same actors do it, or sometimes they sell the access, and the next guy does it. But the key point that I'm wanting to start with is that we have seen ransomware for a while. There's nothing new about this. These are the same crimes. It's just the use of new tools to make those crimes more efficient for the bad guys to do.

In this particular example, at a high level -- and I'm just going to go through step one, two, three here, just to talk about what we saw this particular ransomware initial access brokerage group doing -- they were seeking to obtain a bunch of domains that were aged and had good search engine optimization potential. SEO is what that acronym stands for. That was step one.

Then they would create and host malicious websites with those domains, and those malicious websites would host malware. That malware then would infect any victim who is going online and searching for any term that would return one of those sites in the top whatever bracket. So rather than using phishing emails or anything like that, this particular group was using SEO poisoning.

We would otherwise call this a watering hole attack or any number of other names.

Let's talk about what we saw them use AI tools, LLM, or agentic for, in this case. For the first step, when they are first trying to obtain a bunch of domains with that high SEO potential, my colleague reported that he saw this group using LLM tools to automate the review of domains that are approaching availability on the secondary markets, whether that's by auction or by what have you.

So rather than having to do any sort of manual review, it was more like, "I want to do a campaign targeting this industry, so please review all of those domains that exist on that list, tell me which have the best SEO potential, and which match that industry." It saves them a lot of work. Once done, they would rank them and then use agentic AI in combination with registrars' API keys to capture as many of those domains as possible, as cheaply as possible.

Step two, then they would also use AI tools to create those landing websites to make them look at least decently plausible, probably emulating real websites, and host malware. So the acquisition of that hosting provider is also done by AI. So they're using AI to obtain the sites, to register the domains, to obtain the hosting infrastructure. You can do so rapidly and at a greater scale than you could if any of those steps were manual.

Then, finally, that last step is going to proceed at the same timeline it always did. So a victim's going to go online, they're going to do a

search for whatever term -- maybe you're looking for a caterer in Seville, I don't know -- and then if one of those sites that return on the first page happens to be one that this bad guy set up, if you click on it, you are potentially infected with the malware, which will then lead to a ransomware infection.

What the case agent stressed to me is that what he sees in this is, due to the use of these tools, you're just having that one group be able to create a lot more of these potential malicious sites at a much faster speed than you otherwise would. That was, I feel, a really indicative example of how this is changing the tools that bad guys are using, but the overarching schemes are, to at least our initial visibility, they're very much the same schemes that criminals have been using for a long time. I think it's important to recognize that it's not revolutionary, but it's iterative. And it's still quite dangerous, as we'll get into.

RAM MOHAN

Thanks, Gabriel. That brings us to Theme 1. Before we get there, let's get the next Mentimeter poll up on the screen. While we have that, here is a setup for this question. We're talking about the democratization of attack and mitigation tools. I'm going to ask some questions of some of our panelists here, because we want to talk about how AI is being used both on the abuse side and on the defense side, and where the current reality is and where it may be heading over the next few years.

Let me start with you, Laurin. In your research, what are the major changes you are seeing in how AI is being utilized to supercharge DNS abuse and the defense against it?

LAURIN WEISSINGER

Yes, thank you. First of all, I think we have to go back to a point that both Rowena and Samaneh were making. And this is the question of what is this AI thing? Like, how do we define it? And especially, where does it start? Because we're always talking about traditional automation. If you look at how AI is defined, a lot of that would also fall under AI. So that's one thing we have to think about, that this is not a new phenomenon, that this has happened since two years or something like that. No, we've had AI, computer automation enabled abuse for many, many years. The thing that's changing is probably more the speed, the scale, and especially the accessibility of what is going on here.

So, for example, in the past, if I wanted to sketch up a little app to show someone how it works, I would have to be able to code it myself. That has disappeared. If I have access to an appropriate model/chatbot, I can whip something up myself. Will it be great? I don't know. Will it be full of security issues? Potentially. But I can show what I'm imagining. And if I run something on my local machine to sort some data or something like that, hey, maybe it's enough. So that's one thing where we do see change. Not that, now everyone is doing stuff that wasn't possible before, but it's like, well, now a lot more people can do it. And the people who

already have knowledge are really empowered to use what they know better.

Because if you're a good coder, you can leverage, let's say, AI coding tools to really speed up your process. Because you know what you're doing, you can give very good specs to your chatbot, like Claude or whatever else you're using, and then you can check what it does and refine it. Again, is it an efficient way of coding in terms of energy use globally? That's a different question. We have this background, historically speaking, again. But we have had tools to create domain names automatically, again, for a very long time.

The next issue we have is that while we see data on a very quantitative level, in terms of like, things are changing here, we don't necessarily know what aspect of this change is driven by a specific use of AI, like just being AI-enabled, or AI agents working relatively autonomously. But that is something that our data do not really show. Especially if we're interested in a specific instance of abuse, this name, this site, something like that, then it gets even harder to figure out to what extent which tools were used.

One thing we shouldn't forget is that there is, especially if we go into this direction of profit-oriented cybercrime, there is an industry here where, in addition to AI, there is significant people power behind this, like people who are working these things and doing stuff. And again, here, the intersection of automation and

the uplift that is happening is probably the bigger concern than AI by itself. It's always in an intersection.

So we are expecting obviously more increase in speed, more increase in scale than what we're already seeing, and now more recently, there likely will be an increase in sophistication overall because of the uplift we're seeing from the last generation type stuff. Not necessarily because now, stuff we've never seen before will be happening all the time, but because it's just far more nefarious actors who can leverage just a higher level.

That said, we see in terms of say, vulnerability research, for example, that Claude Mythos -- being the one that everyone talks about -- really is capable of finding stuff that was in long-term open source projects for years or decades in terms of vulnerabilities that human researchers did not find, and now with the aid of these tools, they have been found.

And Claude Mythos is not the only one that does this. A lot of the AI security research firms came in and said, "Hey, you know, we can do this too, and we have for a while." In short, DNS abuse is not easily identified on a case-by-case basis. So this is going to be tricky. There will be a lot of adjustments necessary in terms of what is happening going forward. It's really difficult to say how things will look like in a year from now, two years from now, five years from now.

We have to keep in mind always the big, bad LLM that we're all talking about. But it's not a DNS abuse or anti-abuse tool, it

impacts on an already complex landscape that then has the capacity to change things. And to change things, we don't necessarily need insane sophistication or something that has never been there before. Scale and speed, and reducing the amount of human work that's necessary to get a campaign rolling, is probably enough in other cases.

RAM MOHAN: Thank you, Laurin, for a comprehensive response to that. Let me come to you, Samaneh. Do you think there will be a ChatGPT moment for AI abuse? You know, a moment where the impact of AI significantly changes DNS abuse as we know it. How can we prepare for changes in attacker behavior on the offensive side? And is there an impact for defenders as well?

SAMANEH TAJALIZADEHKHOOB: Thank you, Ram. I would like to echo what Laurin said. I think some of what we are calling AI today has already been in place and been used by researchers and companies for years. The automation and the models. For example, what we partially use today in the large language models to bypass language barriers to translate text comes from a discipline called natural language processing, which has been developed and is being developed right now, as we have the models, because it's still lacking expertise on certain languages.

So, having that said, the DNS I've used as a space, the reason why it's effective is because there are still humans there. For it to exist, we need the weakest link to be able to be lured in a certain way and give the attackers the money.

So as long as the human is in the loop, everything that is developed on the attack side is basically going to be connected to the weaknesses in human languages and skills, language capacity to make decisions. Basically, skills in how internet savvy we are, the social skills that are now social engineered better by the language models, etc. I'm saying that to show that there has not been a lot of innovation in what we call DNS abuse today in terms of attack methods. If we really look at the attack methods, there have always been maybe a few main vectors, and the rest are just playing with human skills around those vectors. If you talk about social engineering, ransomware, and even phishing always more or less stayed the same.

So if there is going to be a change caused by AI on DNS abuse -- and this is an opinion because the short answer to Ram's question is, I don't know, but based on what we have seen so far and the economics of security in general, the change is going to be reactive into how humans are going to change when we are going to play and work more with AI as well. I will give you an example.

At this moment, I think everybody is able to detect if a text is written by AI, right? So, more or less, we are becoming experts of detecting if somebody used AI to write a text. It's fine, but we are able to say

it. So these attackers will evolve on top of this to be more advanced, to be more skillful than us. Unless there will be DNS abuse at one point, which does not include humans, will be machine-to-machine, to which case I really don't know. But as far as it is defined today, I think it will be more reactive to how we develop over time.

RAM MOHAN

Great. So the first shift is about accessibility. More actors can now conduct increasingly sophisticated abuse and defense. The second shift may be even more consequential. Those activities can be automated, scaled, and adapted at speeds which challenge traditional approaches. To cover this, let me hand this off to you, Mason.

MASON COLE

Thank you, Ram. Our next topic is the eventual move toward abuse at machine speed. This is a structural change in DNS abuse dynamics that, in the AI, is "improving" the quality and effectiveness of attacks, all enabling rapid content generation, automated infrastructure creation, and targeting that can outpace traditional detection and mitigation approaches. We're going to talk to our panel about this. I know we're slightly behind on time, so we're going to try to condense some of our questions and answers here.

But first, let me turn to Gabe. Gabe, you gave a real-life example a moment ago in your presentation, but can you expound a bit on what you're seeing in terms of the use of AI and DNS abuse in real life, and how soon are we approaching a point where human-scale defenses can't quite keep up?

GABRIEL ANDREWS

So, addressing human-scale defenses, let me focus very specifically on the law enforcement piece of that for law enforcement as a critical player in society. I think for a very long time, we've had this fundamental assumption that if someone goes out and commits a crime, sure, you don't really have a time machine available to go back and stop it from having occurred, but you can apply consequences to the bad guy when they commit a crime. I think that part of that assumption is that if you capture the bad guy, you can stop them from continuing to cause harm.

And what I fear -- and this is a fear of my own, not from my agency, so if it's boneheaded, throw tomatoes at me, but I fear that we're approaching a point where the availability of automation tools to include AI, whether it's agentid or large language models or what have you, to automate each step of a crime scheme's lifecycle, whether it's the targeting of the victims, identify who it is you're going to send your phishing messages to, whether it's the acquisition of the infrastructure that's needed in order to set up an email provider or hosting provider and so forth, or to craft the phishing messages, whether it's to accept any proceeds of that

crime, you know, Bitcoin payments for ransomware, etc., and then to convert that proceed into the reacquisition of new infrastructure. I don't see any part of that anymore that isn't automatable.

And so I question now what happens when you arrest the bad guy, and the crime doesn't stop. And I think that we're going to be confronted with that real soon. I haven't yet bumped into that, but I think it is just around the corner. And I think that's going to be a profound shift.

MASON COLE: Thank you, Gabe. Let me turn to Rowena next. So, Rowena, you hinted in your presentation about the impact of AI changing the economics of abuse. And many in this room focus on the infrastructure of abuse, but can you talk a bit more about the economic impact on the technology of abuse and what that might mean for the folks in the room here and the people that we all come here to represent at ICANN?

ROWENA SHOO: Sure. Thank you so much, Mason. I think we've had some pretty consistent messaging today around the fact that no one really knows for sure exactly what the impact is here. And I think that runs through to this question as well. In terms of attackers, it certainly seems the possibility of democratizing their access to these tools and making the economic cost of conducting these

activities lower. But it's difficult to have any strong evidence that really gets to that point.

I think when we look at what the defenders can do in this space, that's quite an interesting question. And I would echo what Laurin was saying around this really being an evolution over time, because there have been elements of automation and machine learning introduced in the practices of registries and registrars already. And just to give you some examples, firstly, some of them come from the ccTLD space. A lot of them have been experimenting with machine learning in terms of identifying things that might be risky and combining that with human intervention so that they can reduce the risk of registrations coming in.

If you look at retail registrars, they will have a payment process provider. Many of them will use anti-fraud protection tools available from that payment processor. And those payment processes have been utilizing machine learning and AI in that detection space for quite a while now. And so they're leveraging the use of AI in various parts of that ecosystem.

And I think we can also certainly think about this from a systems perspective. If the scale and the speed are going to increase, how can we bring more tools and services available into a centralized space to make sure we're getting actionable evidence reports to registrars as quickly as possible, and helping build tools and services that facilitate them and help them understand where abuse is concentrated in their zones and how to tackle it?

MASON COLE

Thank you, Rowena. Let me turn to Samaneh now. Samaneh, the bad actors are using AI to lower their costs and increase the attack surface, while defenders are on the other side using AI to increase their defensive costs and their labor requirements. How big is that asymmetry, and how is that going to develop over time?

SAMANEH TAJALIZADEHKHOOB

Thank you, Mason. So, traditionally, the security field is known for having an information asymmetry in a way that attackers always know more than the defenders because they have to attack a specific area, whereas the defenders defend a whole area. On top of that, there have always been barriers to entry across fields. So security experts do not know how to work with AI, and the other way around. So at this point, I would say, and I'm going against some expert ideas that say defenders are behind, I would say maybe for the first time in the history of security, defenders have the same head start to break the barriers that were there before for them specifically. That is, they needed security experts in multiple security areas.

So we needed a DNS abuse expert as defenders, a person who understands network security, application security. At this moment, those barriers are less, so AI helps with that. And also, even for the attacker side, even if it's easier to create phishing text, it's easier to gather personal information on people, it still requires

certain expertise to be able to launch successful attacks on those that are defending with the same tools.

So I would say, looking at what is happening and also looking at the space right now, it cannot be easily said that there isn't a symmetry there. In fact, it could be that they are going hand-to-hand in terms of attack and defense.

MASON COLE: Well, thank you, Samaneh. That's an encouraging word on all this. So let's close off that topic. We've explored how AI is lowering barriers to abuse and still increasing its speed and scale. But the last question that we have to the panelists is probably the one that everyone here is most interested in, and that is, what should the ICANN community be doing about it, and where can coordination make the greatest difference? Ram is going to lead this discussion. Ram?

RAM MOHAN: Thank you, Mason. So we've spent some time here with the panel discussing how AI is changing the threat landscape, lowering barriers to entry, increasing the speed of attacks, and altering the economics of abuse. The question now is not whether these changes are coming. The question is, how does the ICANN community respond? So no single registry, registrar, security company, or government can address this challenge alone. As abuse becomes more automated and more adaptive, resilience will

increasingly depend upon coordination, information sharing, and collective action across our ecosystem.

So that gets us to our final theme and questions for our panelists. What role can ICANN or the community play in helping us prepare for an era of AI-enabled defense and abuse, and where can coordination make the greatest difference? And let me start with you, Gabriel. So, to what extent are we as a community, ICANN.org, defenders, etc., responsible for communicating any changes in how we are defending?

GABRIEL ANDREWS So I think the answer is in the question, that we are all responsible for communicating with one another to share the unique perspectives that we have from our places in the ecosystem. So, for my role as law enforcement, one of the obligations I feel that I have is to show up every so often and talk to you all in the broader ICANN community about what it is that we're hearing from victims that come to us and complain about the crimes that have occurred against them, that cause our investigations to occur, and how that's happening. I think if we don't communicate that, then we're failing a very important obligation.

Similarly, I think it is very important for those that operate these very ecosystem-essential businesses, whether you be a registrar or a registry or a hosting provider or what have you, I think it's very important that you share with us your experiences, whether good

or ill, about how the law enforcement's efforts to engage have either worked or haven't worked and why.

I want to call out, before I close the mic, in my last example, I talked a little bit about my fear that crime is going to soon get to the point where it's automated. And to drill down a little bit deeper, when I, as a law enforcement officer, wanted to, just as an example, get a court order to seize a domain name because it is bad and crime is happening on it now, that process, on average, took me about 30 days.

We're going to get to a point where, if this is automated, you can have bad actors that automate the acquisition of new infrastructure at a far greater speed than I can ever hope to act on it at my scale. Maybe it's faster for you on abuse teams. At that point, we have to really consider, well, how can law enforcement be an effective partner in educating those abuse teams about what's happening so that we're enabling you to act at a scale we can't anymore? And I think these are the kinds of conversations that we're just going to have to have more and more as we encounter the speedier, the greater scale abuse.

RAM MOHAN Thank you. Samaneh, question for you that is kind of more focused on ICANN.org, OCTO. So, how do we ensure that the ICANN.org team, the compliance team, perhaps, or OCTO, takes into account the impact of AI in their role? You know, you look at the compliance team, they have a role of holding registries or registrars

accountable. But how do you do it when a domain is good when they check it, but is evil when a visitor gets to it? And do you think that in this community, we should be looking at AI as a technology trend or actually as a challenge to the resilience of the DNS system that we're here to steward?

SAMANEH TAJALIZADEHKHOOB

Thank you, Ram, for the question. I would say ICANN is not really different to other organizations in terms of tackling this technological evolution. Similar to others, we need awareness, evaluation of this, and application of the tech that we are going through. That said, two more important points. One is that it's really not possible to say something is AI. We cannot say this attack is because of AI or run by AI -- what Laurin and Rowena already pointed out. It's because AI is just a tool that makes it faster or better, but we cannot really attribute if this is caused by AI.

But what we can do, and we are doing as we, OCTO, as the research function of ICANN, and compliance as the function that is looking at DNS abuse for a specific meaning, is to be able to evolve our methods to make sure we consider DNS abuse as it's happening faster. So basically, our measurement methods should reflect that we are doing that, or it should reflect the newer type of domain generational algorithms that are being used.

Also, what compliance has been doing and will be doing more and more is that they are using all kinds of data that they can get their hands on, just to make sure that they are angulating their method

and not missing anything. And this becomes more and more important given the use of AI for DNS abuse. Thank you.

RAM MOHAN Thank you, Samaneh. I'm going to come to you, Rowena, with this question. If AI is changing the threat landscape, the economics, what do you believe ICANN is uniquely positioned to do?

ROWENA SHOO Thanks, Ram. So I'm not sure anything about what we've discussed today changes the fundamental role of ICANN in this ecosystem. I think what we've been talking about is how AI is impacting all businesses and individuals and is having, potentially, a change in the scale and the automation, the speed that we see the abuse created.

So I think it is really important to keep talking about this as a community and talking with each other about what we're seeing and what we're observing, trying to understand what we do know and what we don't know, because we can only really react to what we see as happening. And it's through those conversations with each other that we can identify commonality in terms of problems that might need solutions. That's certainly something that we try to do as the Institute.

If we can see something where we might be able to make a tangible impact at a system level, we'll try to build that out and make it accessible because all of our services are offered for free. So to

continue having those discussions about what the problems and the concerns are as they come up will help us identify tangible courses of action.

RAM MOHAN Thank you. And thank you to all the panelists here for sharing your insight and experience with us today. Now, let's hear from the community. Kathy, you're going to help manage this along with Mason and me. And can you just tell us what we're going to do?

KATHY SCHNITT Sure. If you have a question, you can actually open the Q&A pod in Zoom and type your question. Or if you would like to ask it verbally, just raise your hand in Zoom or go ahead and raise it in the room. If you're at a microphone at the table, when we call upon you, obviously, use that. If you're in the audience, standing up, and you want to ask a question, raise your hand. We come to you. There's a roaming mic over here.

And with that, we'll go ahead and take our first question online from Nico from Internet Society. "If AI enables attackers to register and deploy malicious infrastructure at machine speed, should registries, registrars, hosting providers, and DNS operators deploy AI-based abuse detection capable of acting at the same speed?"

RAM MOHAN Who wants to take that question?

GABRIEL ANDREWS

So I think that it's important that we recognize that as a threat. So yes, I think that machine speed acquisition of internet identifiers is something that we all should be thinking about. I don't know that we need to be so prescriptive as to say that the only possible solution is the utilization of AI on the defender end.

I think here in the ICANN space, one of the things that we're already contemplating might have potential to really be helpful here, and that is we're contemplating a policy on introducing "friction," because I don't really know what that means yet, to the API usage to obtain domains in bulk.

And this is something that maybe actually Samaneh might speak to a little bit, but I know that in one of ICANN and OCTO's past reports, the correlation between the usage of APIs for domain registration had a very high degree of correlation to the potential for future abuse. I think that itself could be quite impactful in addressing one of my greatest fears, and that is the automated acquisition of these identifiers.

MASON COLE

Gabe, thanks for the answer, and thank you for the question. Let's go into the room for just a moment. Does anyone in the room have a question they'd like to raise their hand for? Let's go over here. This is the first hand I saw.

HOUDA CHIHI: Okay, thank you so much. So my name is Houda Chihi. I am an ICANN Fellow 86. So the question is, there is some initiative in building their agents. So how is ICANN dealing with this? Thank you.

MASON COLE: I'm sorry. I think we had a hard time understanding the question. Do you mind repeating it? Thank you.

HOUDA CHIHI: Okay. So the question is, there is some initiative of agents or from companies that are building their own DNS, call it agent DNS. There are at least two, one from the UNIX Foundation and another one is just a draft submitted to the ITF. So, in this context, how will ICANN deal with this?

MASON COLE: Laurin. Please.

LAURIN WEISSINGER: Okay, yes. So I think we'll see on the side of what ICANN does with this. There are multiple proposals, as you're mentioning, that are essentially being pushed into the ITF right now. I think one of them was released in, let's say, March, to essentially leverage the DNS to identify agents. So it's like, I want to find that agent, that functionality, but I would go through the DNS to do so.

Now, on a technical level, DNS can do this, but it is meant to do essentially this type of work. What it would mean on a policy end and how it would be implemented, I think there are a lot of open questions on how one would deal with this. And frankly, I'm not sure, at least I'm not aware of, these discussions having happened in the ICANN space. But it's definitely something to watch, without question.

MASON COLE: Thank you, Laurin. Let's go online next. Kathy?

KATHY SCHNITT: Yes. Our next question is from Kaveh Ranjbar. "I have a question or comment I would like to address to the panel when the time is right. It is about a new category of attacks emerging, as well as a new category of discoveries, looking at the larger system using DNS as one of the vehicles."

RAM MOHAN: So I guess that means, Kaveh, you have a question. You have to unmute and then speak.

KAVEH RANJBAR: Hello, everyone. And thanks for the time. I will try to be brief. So basically, I think my mind is maybe more practical and possibly easier to act on. Which is basically my suggestion, which I would like to read what the panel thinks about it -- is to have an input

channel for ICANN, for people to be able to disclose things which are being found by AI. And I will explain why I think ICANN is needed for that, because there is a category of issues. Let's say a researcher or someone who is doing work with AI comes up with it, and there is no clear disclosure path for any of those.

I encountered one recently, and I will explain how I am going to address that, but I think this is going to happen more. So we have a technology in our company, a graph of the whole internet. What happens is looking at typosquatting, I found what I think, I might be wrong, dormant infrastructure of a full set of payment providers, government organizations, all owned by one or two entities, dormant, pointing to a specific IP address on ASN.

I think that's a setup for possible future attacks. But there is no single entity owning this. This is not a registry or registrar problem. This is not a software vendor problem. And this is not a specific government issue, even. And there is no channel to communicate that. I will just briefly mention what I'm going to do for this type of stuff until there is a channel.

RAM MOHAN Sorry, Kaveh, we are out of time for your questions. Does anybody here want to respond to Kaveh's question?

ROWENA SHOO I think that sounded like a comment to me.

RAM MOHAN Okay. Sorry, Kaveh, your question didn't come through clearly, so we'll have to move on. Let's go to the room. Another question from here in the room. I'm going to go to my right and pick that gentleman there.

ANDREW CAMPLING Hi, thank you. Andrew Campling for the record. Noting the Mentimeter responses, which seemed reasonably positive. When I went to the RSA conference, the view there was that the balance of agentic AI was pushing it to attackers, not defenders. Registries and registrars have to defend against agentic abuse of the DNS and also red team attacks on their systems. So does the panel think that the room is way too optimistic on where their current balance is, and it's really with the attackers at the moment? We're struggling. Thank you.

RAM MOHAN Okay. Who wants to take that? Are we being too optimistic?

ROWENA SHOO Thank you very much for the question. I think it's an interesting perspective. I think what we've presented here today is the best information we can find based on our research and our analysis of this issue. And it certainly seems from the data presented by all of the panelists that we're not seeing a significant change here in the way that you're describing. Of course, we do need to be

considering how we can change and evolve to fight against these attacks.

And there are significant ways that it can increase the acceleration, the scale, and the quality of the targeting and the conducting of these attacks. But perhaps the defensive mechanisms remain the same. We do need to be trying to up our game, but it's still about acting at the DNS level, where we have sufficient evidence, and considering how we can make sure that we're not taking in malicious registrations at the point of registration. Does anyone want to add?

SAMANEH TAJALIZADEHKHOOB: Thank you, Rowena. What I would add to what Rowena said is that indeed, what we discussed today, and also what our data shows, is based on the data collected on how DNS abuse has been measured traditionally. So if the space changes, the data collection changes, then we see different things. For now, this is what we are discussing, what the data shows based on how it's collected. What we will be careful on, and we are not definitely optimistic, is that we are going to keep an eye on if the methods are changing, then we are going to change our measurement methods as well. That's it.

MASON COLE: Thank you, Rowena. Thank you, Samaneh. Kathy, back to you, please, for another question from the pod.

KATHY SCHNITT

Thank you. We have a question from an anonymous participant. "Currently, there seems to be no link impact between mythos and DNS. Do you think this is a link that could be developed in the near future?"

LAURIN WEISSINGER

Okay, I'll take this one. So we have seen that Mythos has been used, as well as quite a few other models, sometimes human-assisted, sometimes not, to look into code bases. And we have seen that these contemporary models, I'll just say, have the capacity to find vulnerabilities pretty much unassisted, and even better, obviously, if assisted by capable humans. So if you now take these models and you look at typical DNS implementations that are being used and related tooling, et cetera, et cetera, it stands to reason to believe that they will also be able to find vulnerabilities and bugs in there as well. So that's coming or potentially coming.

The question is, how do we deal with this? Because a lot of the tools that we're using at this technical end are open source. So, how do we manage kind of the resourcing here, and are we able to respond quickly enough? If you look at what the open source community says about these issues, some are overwhelmed, some actually see a higher quality of reports now compared to the past, where they got a lot of slop. So the question is not just, okay, is it possible? The answer is yes, most likely. The question is, okay, and how do we

deal with it, how do we make it better, and how do we leverage these tools to make our implementations, our code more secure?

MASON COLE: Thank you, Kathy, back to you.

KATHY SCHNITT: Thank you. We actually are going to take one question -- he had one of two, actually -- from Maxim. We're going to take one. "Has ICANN considered contacting leading AI platforms so they can track attacks done using their own platforms?"

SAMANEH TAJALIZADEHKHOOB: We did not particularly consider contacting AI platforms for that reason. What we constantly keep an eye on are platforms that collect good data, which would help us conduct research and know more about this space. If that company happens to use AI as a technology that gives us better data than there then that would be our choice, but not specifically for that reason.

MASON COLE: Thank you for the question. Now, let's go back to the room quickly. Anyone else would like to raise their hand? Over here, please. Vivek.

VIVEK GOYAL

Hi, Vivek Goyal, COO, LdotR, online brand protection company. Today, the way we see DNS abuse, it's directed towards humans. A human has to do something to fall prey to the abuse: lose money, give their credentials, or something. We are increasingly seeing humans rely on agents to execute things, buy things on their behalf, make purchases, and even do money transactions. I would like your views on how abuse will evolve so that it is not targeting humans, but it is targeting agents. So, agents abusing other agents, and what role, if any, ICANN has to play in that. Thank you.

RAM MOHAN

If you don't have an answer, you can just say you don't have an answer. Don't feel compelled to have an answer if you don't have one.

ROWENA SHOO

Sure. I think there's maybe some interesting themes coming up here because I'm thinking about that previous question, which is around, are we being too optimistic? And I think a key differentiation perhaps between the discussions, the question asker was talking about, and the discussions we've had here today is that those discussions were probably thinking about the cybersecurity of the entire ecosystem of everything that uses the internet and all of the networks connected to that, and all of the code involved. I think this panel has focused very tightly on the malicious registration of domain names because that's what's

most within the remit of ICANN, and therefore, I think we saw that as the most relevant here.

But I think there's been some really interesting discussions raised by the last two questions as well that are more going to -- I think one was kind of getting at the trust and safety of the use of chatbots. It was kind of more about what's going on within that chatbot and how they could be abused. And then this question, I think, is a really, really super interesting one. I'm not sure I have a specific answer. I don't know if anyone does, as to what we do when AI is defrauding AI or phishing AI. Perhaps we don't want to give over our wallet credentials to ChatGPT just yet.

MASON COLE Thank you, Rowena. Kathy, over to you.

KATHY SCHNITT Thank you. We have a question from Godsway Kubi. "AI is becoming both a tool for combating DNS abuse and a tool for attackers. How can the ICANN community ensure that AI-driven abuse detection systems are transparent, accountable, and proportionate, particularly in avoiding false positives that may impact legitimate registrants, while still effectively addressing increasingly sophisticated AI-enabled threats?"

MASON COLE Rowena?

ROWENA SHOO

Thank you for the question. I think this comes back to the importance of not just purely relying on AI to do the detection and mitigation. I think it's important to have a combination of optimizing the things that AI is good at in terms of summarizing, adding context, and perhaps collecting evidence, and then combining that with human oversight and human interaction. Because it is an important decision to make if you're going to suspend someone's domain name.

And I think the combination of those two things really helps. Certainly, false positives are a really big concern here. And if you accidentally take a legitimate bank or business offline, that's a really big deal. So I think, actually, the people pressing the button to suspend domain names are probably quite aligned with that concern around making sure that that's the right decision. And part of that will be understanding how the decision is being made if you're incorporating AI into that process.

MASON COLE

Okay. Thank you, Rowena. Kathy, I believe we have a hand up from Theo.

KATHY SCHNITT

Yes, go ahead, Theo.

THEO GEURTS

Yes, thanks. Just an observation, guys. This is really novel talking about AI, but nothing of this is anything new when we talk about automation. I mean, a decade ago, we had the rise of the crime as a service platform for everybody with half a brain cell could pull off phishing attacks, malware attacks, or God knows what. And as a community, we never got any policy work around that coordination or any thinking in that direction.

When we look during the pandemic, we saw the criminals move to human trafficking, so they would have a human pool for half a million people being locked up to pull off all these crypto scams. And now we're talking about automation through language models. But the underlying issue about all of this is not being addressed by any of this. I mean, these language models will just be another form of automation, while we have already seen the other forms of automation, and they're still going on. So I think when you talk about the approach here, we need to do some thinking. Thanks.

MASON COLE

Thanks, Theo. I just want to confirm. That sounded a bit more like a comment than a question. Did we miss a question from you in there?

THEO GEURTS

No, it was an observation.

MASON COLE

Okay. Thank you for sharing. All right, I think we have time for one more in the room. So let's go to Pinky, please.

PINKARD BRAND

Thanks, Mason. Pinky Brand for the record. Thanks to the panel for a very informative discussion here. One of the things that struck me in the conversation was the fact that, you know, does it really matter if it's AI or not? And it strikes me that from an operational and compliance point of view for registries and registrars, that really what the problem might be is more about a KYC issue and the trust signals that registry operators and registrars need to have to know not is this real or not, but what trust signals are there that tells us that this is actually real and that we're not taking down something that shouldn't be taken down.

So my question is, knowing that some registry operators have proposed some sort of KYC, know your customer verifying agents out there -- I believe there's one from GoDaddy and another one from Identity Digital, both very interesting proposals out there -- does it seem to anyone that there should be one common standard, perhaps, or that there needs to be more coordination about KYC and the rise of agentic, these agents out there? Because there will be registrars in the future that are completely automated and not even dealing with humans. Thank you.

ROWENA SHOO

Thanks so much for the question. So it's interesting, I think of KYC as potentially one element that we would associate with trust, but it doesn't have to be the only element. And I think there are other parts that registrars could look to in their systems to understand if they can trust the customer at hand. So that could be whether they are a longstanding customer, they've renewed a domain, or whether or not they've had previous complaints against them.

I think standardization is tricky here when you're looking at the entire world and different legal systems, different capacities, and ways of operating. And also just to be a little bit careful around whether we're creating a new market for false identity, stolen credentials.

LAURIN WEISSINGER

I think there are also other issues, obviously, to consider. We, again, have a lot of technology here that's actually useful. Various countries in Europe allow you to use your ID card to actually prove you're real without necessarily sharing all the personal data that's attached to it. At least where I grew up, while the capability is there, it's hardly ever used. And the way it's often done, it's very data-intensive. Lots of privacy issues that I think Rowena was also referring to.

So trust signal's great, but the question is, how do we do this properly in terms of like protecting people's data? Because we know, for example, how often breaches happen. And then the next step is, as you're saying, if it's an agent acting on behalf of another

agent or a person, how do we deal with them, how do we identify them, how do we build these traces? And that's it's actually very complex, especially internationally.

MASON COLE: Samaneh?

SAMANEH TAJALIZADEHKHOOB: Thank you, and I would add to that that maybe more than ever now, because of the lower barrier to entry for attackers, anything that we come up with, let's say KYC is an example of a trust signal, the attackers will adapt to that quicker than before. So it's more dynamic.

My research team is the one behind the work of associated domain work, which ended up being in the PDP process. And the signals we initially identified as signals for associated domains are already becoming weaker signals because attackers are adapting according to those. So no matter what we choose as a framework, it's going to change faster than before. So that element is more important than before. Thank you.

RAM MOHAN: Thank you so much. So, Gabe, 60 seconds. What do you believe needs to happen first inside the ICANN community on this topic? There's a timer up there for you, by the way.

GABRIEL ANDREWS

Sure. I'll go even faster than that. Piggybacking on what we just heard from Samaneh, I think that the upcoming policy development process that's going to deal with access to the API keys that enable bulk domain registrations is going to be a very important thing that we can turn our attention to. That is very much something that ICANN can address.

Potentially, even the prior question about KYC, because I know it is cost-prohibitive -- and I've heard this many times from our registrar partners that KYC costs money -- perhaps it is more relevant and applicable to that much narrower set of customers that require access to APIs enabling bulk domain registrations. So these are the kinds of things that we can contemplate without being wed to any ultimate outcome.

RAM MOHAN

Rowena, what do you believe needs to happen first inside the ICANN ecosystem? Sixty seconds.

ROWENA SHOO

Thanks, Ram. So firstly, continue the great work we've already started on the policy development processes. Secondly, expand our research and our understanding of the evolving threats. And I think actually, OCTO is doing an excellent job of that. And then finally, share our experiences and our understanding with each other so that we can understand what's coming next and evolve and change.

RAM MOHAN Thank you. Laurin, same question. What do you believe needs to happen first inside the ICANN community on this topic?

LAURIN WEISSINGER Very similar, if not the same answer to the same question. I think we really need to work on our understanding, and figure out what are the possible future scenarios, the key drivers here, and then figure out how do we respond. Especially, one of the questions before was about what are the agent identifiers in the DNS? These are questions that will come up, and the community really needs to figure this out.

So, for example, we in SSAC are working on this. There is, by the way, a work session on this tomorrow, if you're interested in the topic. You can see if we have room in, I think, a space that's not as big. But this is really where we have to go. We have to figure out, okay, what might be happening and how could we respond?

And then, obviously, policy needs to be created. And one of the big challenges will be the relatively high speed of how the space is changing versus how long it takes to deliberate and make policy. This is a general issue, not just ICANN. Standards, and then all these things also run behind what is technically possible. And, you know, you can only try to catch up as fast as you can.

RAM MOHAN Thank you. Samaneh, you have the last word on this question.

SAMANEH TAJALIZADEHKHOOB Okay, I would say more than before today, with the rise of information and information about everything, the most important thing is to remain more critical, and critical thinking for us as a community, and also keep an eye on the work on this topic that is being done in other communities. So, more importance of cross-community work.

For example, I know there is a mock community who work on this topic of DNS abuse and domain names. There is a community in the IRI, and a lot of work in the RIR communities that work on these topics too. So keep an eye out because there will be more use of cross-technologies than before. Thank you.

RAM MOHAN Okay, thank you. Over to you, Mason, to wrap this session up.

MASON COLE Thank you, Ram. Well, we tried to get to as many questions as we could. Sorry, we couldn't cover just quite everybody. But as mentioned, this session is informative. It's not prescriptive. So the time for policy development is going to come soon enough.

We have a lot to take away from this session, including, for example, we're on the cusp of AI-influenced abuse. We need to be on guard for that. The economics of abuse are changing, as well as

the instrumentalities of abuse. And there's an opportunity not only for further impactful harm, but also an opportunity to help defenders.

So, on behalf of the BC and on behalf of the SSAC, thank you to our panelists. Thank you also to Carlos and Kathy of ICANN.org for their help in putting this session together. Thank you both. And especially thanks to Ram and the SSAC for co-sponsoring this session.

This has been a particularly informative use of time on a developing topic that we, as stewards of the domain name system, are sure to have to address, and likely very soon. We thank you all for being here. Thank you very much.

RAM MOHAN Thank you, folks.

[END OF TRANSCRIPTION]