ICANN79 | CF – At-Large Plenary hosted by NARALO
Tuesday, March 5, 2024 – 4:15 to 5:30 SJU

MICHELLE DESMYTER: Thank you. Hello and welcome to the NARALO Roundtable, the At-Large Plenary hosted by NARALO. My name is Michelle DeSmyter and I'm the participation manager for this session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior. During this session, questions or comments will only be read aloud if submitted within the Q&A pod located at the bottom of your Zoom screen.

If you would like to speak during the session, please raise your hand in Zoom. When called upon, virtual participants will be given permission to unmute in Zoom. On-site participants will use a physical microphone to speak. Interpretation for this session will include English, French, and Spanish. Please state your name for the record and the language you will speak, if speaking a language other than English. Please speak at a reasonable pace. And with that, I will hand the floor over to Greg Shatan, NARALO Chair.

GREG SHATAN: Thank you very much and welcome this afternoon to the NARALO Roundtable. As many of you know, it's the At-Large tradition that the NARALO, who is the host RALO, who's part of the world we are in, not only acts as host, but acts as host for one or two panel discussions and the like. This is our moment to discuss an issue of great current interest

and has been of interest, of course, in some ways since the dawn of the web or even the dawn of the internet in some ways before that, although at least it wasn't DNS that was being abused at that time. We have a distinguished panel with us, or at least most of them are with us. In order from my right, we have Laureen Kapin, who is the Assistant Director for International Consumer Protection in the Office of International Affairs in the U.S. Federal Trade Commission and also outgoing Chair of the GAC Public Safety Working Group.

I want to congratulate her and thank her for her years in that role as well. We have next to me Graeme Bunton, the Executive Director of the DNS Abuse Institute and an actual Canadian. To my left, Steve Crocker, who in many ways needs no introduction. You can read about him in many places, but it is really both an honor and a pleasure to have him in this panel. Many things I admire about Steve is he never stops learning and seeking and taking on issues. It is great to have his focus on the issues of DNS Abuse now. There is nobody, in a sense, with more credibility and gravitas than Steve. I can even feel the specific gravity coming off of him. Steve is CEO and President of Edge moor Research Institute Incorporated. Next to him, we have Ram Mohan, who has worn many hats and is currently the Chief Strategy Officer of Identity Digital.

Identity Digital is what happens when you put a donut into a foliar vice versa. We are pleased to have Ram with us. Then we have Reg Levy, who can come up to the front if she wants. She just sat down. I am making her stand up again. She has to walk around with her coffee and her mask on and her computer and all sorts of stuff. Reg Levy is the Head of Compliance for Tucows. Reg is also someone who has toiled for years in the garden of ICANN. We are glad to have this distinguished

panel to discuss DNS and domain Abuse in the digital economy. Since this is a panel from the At-Large perspective, the end-user perspective, we hope that the end-user perspective will be duly taken into account in the conversation.

There have been many conversations at this ICANN meeting, as in previous ones about DNS Abuse and several interesting developments that have come to fruition fairly recently, although they have been in the works for a while. One is RDRS. Another is what are being called the Domain Abuse Amendments to the Contracted Parties Contract. We also have a variety of issues. We also have after many years, at least a settled working definition of DNS Abuse that at least people agree is the minimum amount to be discussed.

We do not have to discuss the shape of the table for years, as we used to. While we have an order here, and Steve even has a capital C next to his name, I'm actually going to start with Graeme Bunton. As the Executive Director of the DNS Abuse Institute, since his entire professional life is devoted to DNS Abuse, I would like to get your thoughts. You can kind of kick us off. You're welcome.

GRAEME BUNTON: Thanks, Greg. That's a very open-ended start. This is Graeme Bunton for the transcript, as introduced. I think the key thing to really-- let's maybe start with the amendments, because it feels like sort of a momentous thing that's happened very recently. They come into effect on April 5th, I think, although I look towards Reg at that. They create clear or clearer obligations for registries and registrars as they relate to malicious registrations.

They give compliance the tools to enforce. The future, I think, is very bright for the reduction of DNS Abuse, contingent on ICANN compliance using the tools that they have now been given. I am optimistic that they are going to use those tools, and we'll see that probably relatively shortly after these amendments come into effect. And I think it's going to change the landscape of Abuse. I think that's going to be very interesting, and I think we as a community need to pay very close attention to how the landscape changes. It's going to inform future work from this community, and it's going to change in many ways how the Internet looks and works, I think.

Maybe this is a spicy thought to put out there, but we're not solving cybercrime. We're not going to make bad people stop doing bad things. What we're going to do is make it more difficult to do that in the traditional DNS and the bits that are regulated by ICANN. How does that change how bad actors work? Do they move to ccTLDs that don't have the impact of these amendments? Do they move to more compromised websites, knowing that they might have longer to exploit those services? Do they move to subdomains which are outside of ICANN's purview? Do they move to alternate routes? That sort of second-order effect of these amendments, I think, is going to be very interesting to pay attention to as a community. Maybe that's enough spicy thoughts to get us started, but I'm really optimistic that we've made some very positive changes within this community in recent history.

GREG SHATAN: Thank you, Graeme. I think that gets us started off very well. Sorry to hear that we're not completely solving cybercrime, but one of the

constants is the truism that the bad guys are always a half-step ahead of the good guys, which is why the good guys have to keep moving, because you don't want the bad guys to be a whole step ahead. Talking about good guys and those who are in the compliance business, I think that naturally brings us to Laureen Kapin as our next speaker. I know Laureen has some slides, and I'm probably causing a lot of, oh, there I go. Nicely done. Okay. I was worried that by going out of order, I was going to screw things up, but ICANN has the best support staff in the world.

LAUREEN KAPIN: Yes.

GREG SHATAN: Go ahead, Laureen.

LAUREEN KAPIN: You did have an order that you were diverting from, but we are nothing, Graeme and I and the rest, if not nimble. I'm Laureen Kapin. I'm speaking in my capacity as both co-chair of the Public Safety Working Group, at least through the end of this meeting, and also as the Assistant Director for International Consumer Protection at the Federal Trade Commission. These are my remarks. They don't represent the views of my agency. I have to say that. You'll see in the next slide, a little post-it tab that says just that.

There you go. But what I wanted to share with you to sort of give a real-world perspective of how certain types of DNS Abuse may play out in

the real world of fraud are some statistics. These are FTC 2023 statistics about the fraud complaints we receive, just by way of background and also for clarity. These fraud statistics are not equivalent to DNS Abuse, which has a very narrow, specific definition in our ICANN world. These are fraud statistics, generally, that the Federal Trade Commission collects because we are the United States National Consumer Protection Civil Law Enforcement Agency. We prosecute false and deceptive practices, and we collect millions of complaints from around the United States, from law enforcement, from consumers and victims themselves, from the better business bureaus, and even from some international entities, law enforcement entities, and private entities. We have a very rich source of data that gives us a lot of data to be able to share with the public. Just a little word about our name.

Sometimes people think, Federal Trade Commission, oh, you must handle international trade. It's trade in the sense of commercial and consumer transactions, not imports, exports, or tariffs. I always do like to say that for folks who aren't familiar with us. This is sort of a highlight, and again, to give a real-world perspective into how certain types of DNS Abuse may play out for the end user. That would be you and I, who transact business on the Internet. Just generally speaking, fraud reports are going up. They're not going down. We received 2.6 million fraud reports last year. That is up. In terms of losses, and this is just in terms of dollars that's separate and apart from emotional anguish and other things that people experience when they're victims of fraud, $10 billion reported lost.

In terms of our top frauds, and this is probably of interest, especially to this community, imposter frauds, it is the top category. That includes

business imposters and government imposters. Of course, we know that one of the topics that's talked about at ICANN are domains which are pretending to be legitimate businesses, but which aren't. Imposter scams is one of those top categories. In terms of contact methods, one of the most popular, I'll say, for want of a better word, the most popular contact method is email. Again, that is a method that is of particular focus in the ICANN community. Just in terms of besides dollar losses, one thing which we also note in our statistics is that people are losing a lot more money to investment scams. When we're thinking about exactly what the subject areas are, investment scams, particularly about crypto currency, are becoming more and more popular.

Second highest reported dollar amount, that's imposter scams, which, as you may recall, is our top category of complaints. This is the dollar loss, though. Again, this underscores that imposter scams are surging, but you may be wondering about other topics. That's online shopping, prizes, sweepstakes, and lotteries. That's the famous you've won scam, investment scams, and business job opportunities. This is a deep dive into contact methods. I already mentioned that email is the most popular, but phone calls and text messages quickly follow. If you're thinking about what are the methods that result in the biggest dollar loss, social media, but then more pertinent for ICANN online ads, pop-ups, and websites.

Then I want to focus particularly on phishing, which is one of the specific types of domain name Abuse that we're focusing on. In terms of our fraud data within the past year, and that's from March 3rd, 2024, we received over 1,000 reports, nearly 100 in the last month, that talk about, complain about phishing. Those reports, in terms of total

amounts paid to scams that may involve phishing, that's 2.2 million. Again, we see that the top product or service that that relates to, that's business imposters, another topic that is widely discussed. I'm going to stop there in terms of our data, but if we can go to the last slide, I just want for your reference, if you want to do a deep dive into the FTC's publicly available data, I would encourage you to look at our data book. I would also encourage you to look at our publicly available data, because you can compare complaints and trends year by year, quarter by quarter. You can see what age groups are most likely to be ripped off.

You can see who pays the most. You can compare contact methods. You can slice it and dice it in so many different ways in real time. It's a very cool resource. I provide these insights because what I want to underscore is that sometimes we get caught up in debates about, well, how should we define this and who you should report to? Those are all very, very important topics to grapple over. But I also think, we can't lose sight of the impact on the end user, which are the real victims here who are defrauded in certain instances of their life savings, their retirement, they may lose their homes.

Some people even sacrifice their lives because they're so distraught. I know that sounds over the top and very dramatic. I do that to point to the full spectrum of harms here, which can range from no loss at all, they're just reporting complaints because they know it's a rip off, to someone who really is very significantly harmed. I think that's important to keep in mind. I'm going to stop there.

GREG SHATAN: Thank you, Laureen. That's some very sobering statistics. I'd like to turn next to the only man who has a form of computer memory named after him, Ram Mohan. Ram, I would like your thoughts anywhere in the discussion area that we have, even with our very narrow technical definition of DNS Abuse, you have a lot to talk about.

RAM MOHAN: Sure. Thank you, Greg. I appreciate it. I want to focus a little bit on actual on-the-ground things that get done, because one of the issues that we ought to be looking at is not the incidence of Abuse and not even the mitigation of Abuse, but the time to mitigation of Abuse. That really is, in my world, probably the most critical, most essential metric to look at. If you look at phishing as an example, which is a scourge that affects all of us in our everyday lives, the difference in impact, negative impact, from taking down or addressing a phishing campaign, the difference between taking it down in the first six hours of knowing about it, to the first 16 hours of knowing about it, merely a 10-hour difference, is an order of magnitude.

So, literally, with things like that, these are practical implications. The focus has to be, I think, on not just how many times something happened, what is the frequency of occurrence. It is, I think, the focus also should not be nearly as much on the volume or the mean time to take down. But we really ought to be looking, as an industry and as a community group, we ought to be looking at median time to take down. And that we ought to relentlessly focus on bringing down. And the way to make that happen, unfortunately, inside of the ICANN realm, a lot of that ends up beating up on registries, beating up on registrars and

saying, if only you did more, if only you did extra. But here's the other part of DNS Abuse, which is that Abuse often requires context.

If you look at something, certain kinds of things may look like they're abusive from a content point of view, but in context are actually not. So, if you deployed automated algorithms that looked at, say, certain phrases that sound like they are hate speech or certain phrases that look like they are focused on criminal acts, unintended consequences of those kinds of automated measures are that research papers get taken down. Or that researchers who are running sandboxes and, in some cases, traps get taken down. So, context, so that's the second point that I'm making. When you're looking at DNS Abuse, context is, I posit, almost as important as the content itself.

The last point is that we look at, we often in this ecosystem, we look at the contracted parties as the people to effect positive change on Abuse and to help make our ecosystem, relatively speaking, safer. But there is an entire set of other players in here who I think we need to find ways to reach out and understand what they're doing. You look at hosting companies, you look at other providers like that. Sometimes you will find that there is something that looks clearly like Abuse. The request comes to a registry and the registry says, in fact, the request may come to ICANN and ICANN says, go to the registry, goes to the registry, and the registry says, go to the registrar. And the registrar says, actually go to the web hosting provider, because this is a content-oriented issue. And we need to find a way, I think, to include them in the solution set.

I think, by the early design of the parties involved in ICANN, those sets of providers are just often not at the table, not here. They're also doing

good stuff. There are many of them who are doing good things, but their best practices are not baked into the things that we're doing. The final point that I'd like to make is that, there is, as Graeme was saying, the goal is not to prevent. The goal is, I think, to mitigate, number one. And number two, inside of our ecosystem, if there are explicit and clear things that we can do to respond more rapidly, to persistently shut down bad behavior, we ought to look at that. Inside of our own registry, we have an experimental program. This is not something that we've done in a production type of environment, but we're testing an experimental program where when domain names come in and get registered, at the time of registration, and in some cases, the registry actually has who the real registrant is.

In many cases, 60, 70, maybe even larger percent of the time, it's hidden behind a proxy. But there are other ways to look at patterns and signatures of who is applying, what kind of strings are applying for. In some ways, the good thing is that criminals also get lazy. And when they find a particular mechanism that is working, they end up doubling down on it. And so, the onus is on us as responsible parties to have programs that identify those kinds of lazy patterns and do something about it.

So, we're looking at an experimental thing inside of our own company where, at the time of registration, if we can discern that there's a possibility that a name or a group of names that are being registered may have not intent, but a pattern that is similar to patterns exhibited earlier that had clear malicious intent, then we put those names on a watch list. And when and if those names actually get to being identified as abusive, the time to take down, we are looking to see if we can get

them taken down in an hour's time instead of six- or seven-hours' time, which is where we are.

And as I was saying at the very start, the order of magnitude impact is dramatically higher the quicker you can shut these names down. You just make them useless for the criminals who are registering them. So, I'm focusing on practical things that can be done. And I'm also advocating for bringing those who are not at the table inside of the ICANN party, bring them to the party as well.

GREG SHATAN: Thank you, Ram. A lot to think about. And certainly, we have talked about how to bring those other parties to the table, and yet they're not here. But maybe that will change. I'd like to turn to Steve Crocker next, because I think that Steve has some interesting things that will segue nicely from what Ram has to say. And even if they don't, they're going to be interesting.

STEVE CROCKER: Thank you.

MICHELLE DESMYTER: One moment, Steve.

STEVE CROCKER: Slides? You don't have the slides?

GREG SHATAN: They should be coming up in a second.

STEVE CROCKER: Okay, thanks. I've been working in an area that I would say is adjacent to the DNS Abuse issues per se, but there's areas of overlap and similarity that are going to come into increasing play. The area I've been working in is what I sometimes refer to as the WHOIS mess of access to the registration data. Actually, both what the collection of registration data ought to look like, what data gets collected, and how should it be handled, and then who should get access to it. There's a project that we call Project Jake after Jake Feiler, who was the original WHOIS person, if you will, way back at the beginnings of the ARPANET. They're coming or not. They're flashing up and down there.

GREG SHATAN: Sorry for the technical difficulties.

MICHELLE DESMYTER: One moment.

STEVE CROCKER: Trying to stay in sync. So, in brief, I've watched over a very, very long period of time. Not measured in years, measured in decades. Repeated attempts at trying to address the questions of what data should be collected, how should it be handled and so forth. And then, particularly within the ICANN structure, watched repeated attempts at trying to settle this out. When I left the Board toward the end of 2017, having

failed to stimulate what I thought was a helpful course, I and several other people went off to collaborate on the side and think hard about the problem. And have put together a set of ideas.

We are now working hard to socialize them and to bring them into play. And at the same time, I've been engaged in the internal discussions, the policy development processes related to phase one and phase two of the expedited. It's hard to say that word, expedited policy development process and also with respect to the registration data request service. A question was asked, what would success look like? And the ideal would be that you have a system that achieves privacy, accuracy and efficiency.

Next slide. Well, try the next slide. Try no slide. Yes, thank you. So, the question is, can you have all three? Next slide. That's the challenge. And the answer is yes, you really can. But it takes some work. And the work that has been done to date has been fairly shallow in the sense of, well, we have to have long discussions over the simple things and we'll never get around to the hard things. And so, the hard things never get addressed.

And GDPR was a big wrecking ball that came and pushed everything over to one side, but it didn't really change the level of detail. And it is sort of equally crude structure, heavily oriented toward privacy, but not very delicate with respect to making nuanced decisions. So, as I said, we've been taking a fresh look. I'm not going to go into great detail on all of this here. I talk about it in a lot of other contexts and happy to talk, but the slides will be available. Next slide, please. So, there's a picture of what domain name registrations look like, simplified in the sense of

you have a registry at the top and a registrar underneath and databases associated with them.

And the data comes into the registrar from, this has been mangled in production here, from the registrant on the lower right. And on the upper left, you have people making requests for that data. And so, this is a picture that was relevant prior to GDPR.

Next slide, please. Let's see. Let's just keep going. This is all legacy WHOIS. Keep going, please. Yeah. So, and the numbers that I understood is that there were on the order of 5 billion queries per month prior to GDPR across the entire system. That's a big number. And when you ask sort of the numbers now, they are, well, they're divided up into the public information, which isn't counted at all. And then the queries for non-public information, which are, well, we get two or three, or we get some number. And the numbers are so completely different that it begs the question of, well, what were all those queries before? And there's a variety of answers about that. Some of which, well, they didn't matter. They didn't count. They weren't needed.

They were just noise and so forth. But I have trouble accounting for multiple orders of magnitude with answers like that. Two or three orders of magnitude maybe, but twice that many is a little hard. Next slide. So, we've been building a comprehensive model. One side of that model is what should be the rules about what data is collected and how it's labeled in terms of sensitivity, what level of accuracy each of the data elements is subjected to, et cetera. And we recognize that when we talk about what the rules would be for such a thing, that everybody

has an opinion. In fact, they have more than an opinion. Sometimes they have actual strong rules. So, a registrar might have rules.

A registry might have some rules. Policy authorities like ICANN or governments will have rules. And these things interact with each other. Some are more explicit and documented in clearer ways than others. But when you tear them all apart, you say, okay, a registrar that is operating might have certain desires about what it wants to do, but it is beholden to the registry. And in turn, both of those are beholden to the various policy authorities that they might be operating under. And there is no a priori guarantee that all of these are consistent. One can easily imagine a registrar that is trying to satisfy, let's say, European EU rules and U.S. rules, and that might discover that in some instances you can't satisfy both of them because they're imposing inconsistent requirements. We've built some tools that make it possible to document what all of those rules are.

Not that we have anything to do with specifying what they should be, but we can document what they are or what they purport to be, and then compare them in a very clean, vivid way to show whether or not they're consistent or not consistent. Next slide. And then on the request side, equally need to tease out the nuances about who gets to select what data and under what circumstances. And as I say, this isn't the right time to get into a lot of those details, but again, a lot of precision is possible if one goes at it from that point of view. We worked, for example, very, very comfortably with Lorraine and Gabe and others in the Public Safety Working Group and got some remarkably straightforward, crisp responses.

We said, so what kind of queries would you like to be able to make and under what circumstances? And it's a gorgeous little picture, which I'll just leave there as a teaser. So, we have request templates. Let's just keep flipping through those quickly. Okay, so maybe go back there, Juan. So, here's a really critical thing. We're sitting in an era where we're told that every single request for data for nonpublic data needs to be examined for balance and for context and for legitimacy of purpose and so forth. And that requires a qualified attorney because the risks are so enormous with every single request.

That doesn't make any sense at all. If you say, oh, well, really, what is it that, that person is looking at each and every time? What data do they need in order to have to make that decision? Why couldn't that data have been collected and be in the system a priori? You're not going to go and ask the registrant typically in real time to give back, to answer more questions.

So, I believe that almost all of the queries that are going to be made and all of the decisions that are going to be made, can be made into a fairly routine and sensible and defensible organization. And by doing that, you accomplish multiple things. One is you reduce the cost enormously. Another is you reduce the time it takes to do that enormously. And the third is you increase the consistency and the transparency of the process.

And the net of all of that is that you bring the process from the 19th century into the 21st century, where it should have been to start with. And some of that involves laying out what the rules of the game are, so that if you say, we'll just pick one of the obvious examples. Intellectual

property attorneys are going to be making some number of requests. You say, all right, if you're going to make a request on behalf of your clients, then here's what the rules are going to be in terms of what your relationship is with them, what their problem is, and you are at risk if you lie about these things. And we know where to find you. And so, if you sign up there and you're respectable and we can audit, then you get to proceed more quickly and we, speaking as a registrar perhaps, can proceed more quickly and things move along at good speed.

And similar kinds of agreements from other groupings. So, this just repeats what I was saying, that in principle, each request is checked, and in practice, much of the above can be automated. Let me break off from that and relate back to the DNS Abuse situation and others. A great deal of these kinds of transactions that we're imagining share a lot with the takedown process for DNS Abuse. Who's making the request, what information is known about that, what documentation, who's at risk if they Abuse the situation, and all of that. And so, there's a bunch of shared, potentially shared machinery that underlies all of this. So, that's my sort of main takeaway in this setting on this topic. I'll just break off here in the interest of time. Thanks.

GREG SHATAN: Thank you, Steve. Well, one of the things I heard from both you and Ram was the concern about time, time to mitigation, time to information. And I think that's a critical aspect of all of this. And if I recall, we were kind of going toward the creation of a grand machine for replacing who is, if you will, and now we have a somewhat less grand machine, the RDRS. Steve, you were donning the hat briefly of a

registrar, but we have someone here who wears that hat on a full-time basis along with a black mask. Take nothing from that. It's just health and hygiene. So, I'd like to turn to Reg Levy from Tucows. Please go ahead, Reg.

REG LEVY: Thank you. And thank you so much for letting me go last because it has helped me address some of the questions that people have raised or some of the concerns that people have raised. I really appreciate Ram's indication that there's a need to take things down quickly. Absolutely agree with that. But we do also have to balance the damage to registrants to the damage that is being done when a maliciously registered domain or, excuse me, not maliciously registered, a compromised domain remains up.

So, I would say about 50% of our Abuse is compromised versus malicious. So, and that's anecdotal, I don't actually have the data for that on me right now, but this is something that we need to work directly with the registrants and directly with the resellers in order to effectuate a change. How easy is it for us to get a hold of them? Do they even know who we are? Like, are they going to believe us when we say there's something wrong with your website? Like, there's a whole host of things that go into that and if we just simply suspend a domain, we're taking businesses offline. And at Tucows, internally and in partnership with other vendors and people who are working in this area, we are looking at ways of looking at actual patterns.

There's been a lot of discussion about registrant data for some reason because I thought we were talking about DNS Abuse. And a lot of times

the nexus is much more closely correlated to an IP address. A malicious actor isn't going to register 20 domains at Tucows because we're going to look at that account once one gets reported and say, wait a minute, these all look kind of fishy and take them all offline. What they're going to do is they're going to register maybe two with us and then two with someone else and two somewhere else. They may use the same often stolen data.

So, it's real data of real people, but not the person who's behind the actual domain. They're not registering Mickey-Mouse, they're registering Reg Levy from when there was a breach of some website that I put my information into long ago. So, the IP address information is a much better indication of an actual ongoing malicious attack because that's where the malicious content is being hosted and it's where it's being distributed from. I said we're working on that internally as well as externally. Tucows has, among other things, a small fiber ISP and we have been looking into ways of synchronizing the information that the domains business gets with information that we can get from the ISP traffic in terms of that IP address information so that we can then go out and say, hey, GoDaddy, just to pick on someone.

Tucows, sorry, a fictional other registrar. Hey, we saw this particular attack on our network. And here's some additional data that we're seeing that correlates it to a domain name that is registered on your system. Can you please take a look at that domain and see whether or not action can be taken on that. Obviously, that's in the future, but this is something that we're looking into because we have found that there is no nexus between registration data and DNS Abuse. So, one of the

questions that we were asked to consider is how can RDRS help reduce DNS Abuse.

Well, RDRS is intended to provide the Board with information about how popular a centralized registration data request system is going to be. That's completely outside the scope of RDRS and the types of requests that we're receiving through RDRS are typically not related to DNS Abuse so much as they are related to content Abuse and that is a completely separate situation that is outside the scope of this conversation, absolutely, and certainly outside the picket fence. Let me see what my other notes said. Oh, Steve indicated that it would be really great if requesters were told exactly what information was necessary so that they could provide it in a standardized format, and then we could easily review it.

And that's absolutely what the RDRS has done, again, not related to DNS Abuse, and we're still working through education with regard to actually providing all the information that we have said that we need so that we can review each one separately because that is necessary. It is not simply the case that because as a lawyer, I specialized in intellectual property, that I get to have access to the entire WHOIS database. Like, I don't understand how that is even an argument. The fact that someone has a job doesn't mean that they are necessarily going to have the right to have access to any particular specific domain names, WHOIS data. Again, I feel like I'm going all over the place here because this is a conversation about DNS Abuse, and I'm talking about WHOIS data, and those are two things that are completely separate from each other.

We've seen time and time again .CN has one of the most strict registration data policies in the industry, and it has an extremely high Abuse level. The Abuse level in .com is not high because it's a thin registry. It's high because it's .com. Everybody wants .com. If I send out a completely legitimate Tucows.link email, people are going to think that that's a phish. That's why .com has so much Abuse because it is still the gold standard. I think those are the main pieces. Oh, no, I appreciate that Steve listed privacy as something that was necessary to this ecosystem because, as Laureen mentioned, there's so much fraud on the internet, and so much of it has come from the fact that WHOIS information used to be public.

That was just an incredible source of free information about people, email addresses, phone numbers, home addresses, likely, that was used to effectuate Abuse against domain name registrants who are end users of other products. These are also people who deserve to have protection. Again, there was no change to the DNS Abuse landscape to how my team takes down domains when GDPR came into effect. They're completely separate conversations.

GREG SHATAN: Thank you, Reg. You've given us a lot to think about, and you had the privilege of going last. But I'll open this up first just to anybody on the panel who wants to put up their hand to either respond to anything that anybody else said or add something new to the conversation. And if nobody has a question or an answer or a statement, I will cede the pod, but I see Lorene has her hand up. So, Laureen, please go ahead.

LAUREEN KAPIN: Sure. Since I gave the parade of horrible statistics, I also wanted to balance that out with some appreciations and picking up on some of the threads that have been discussed. First of all, I think the DNS Abuse amendments are a great development, especially as they were initiated by the contracted parties and ICANN. I think that sort of proactive initiative to give compliance the tools it needs, especially to deal with recidivist or systemic bad actors, is a really welcome development. Other welcome developments are the great metrics we're getting reported by our colleagues at the DNS Abuse Institute.

That sort of visibility into what's going on can help inform and drive future policy efforts, and I think that's great. I know ICANN. org is also moving forward with its Metrical system, and I'm really looking forward to seeing what that looks like, because it sounds like that will be more granular information at the registry and registrar level. I also was encouraged to hear about the proactive work that registries and registrars are doing to try and detect potential bad actors and actresses at the outset, instead of just reacting to a complaint. I think that sort of work is tremendously important and could really serve as a best practice and a model for the community to be very energetic and enthusiastic about trying to deal with things in advance, as opposed to just reacting to a bad situation or, worse yet, a crisis situation. I did want to add those appreciations.

GREG SHATAN: Thank you, Laureen. Reg, I see your hand up, but first I'm going to push back perhaps a little on the idea that there is no connection between DNS Abuse and registrant data. In thinking about this, we're talking

about both bad acts and bad actors, and it seems to me, DNS Abuse itself is a bad act, and there's no connection directly between WHOIS data or RDRS data and the bad act itself.

But when you're trying to look for the bad actor who did the bad act, then either who is data or some other data or some other way, and as you mentioned IP addresses, and following those breadcrumbs, but whatever they may be, will lead you to the bad actors. And in order to stop more bad acts, among other things, you try to find bad actors, you also try to wall off their methods and their supply, et cetera, et cetera. But that's just a thought that I'm having, and your hand is up entirely independently of me saying that, so I will go to you and then to Graeme and then to Steve.

REG LEVY: Thank you, and thank you, Lorraine, for raising the point that I had intended to make in my introduction, that the DNS Abuse amendments are intended to give ICANN compliance the teeth and the tools that it said that it needed to go after the bad actors. And when I say that, I mean the registries and the registrars that actively support and profit from DNS Abuse and from failing to do anything about it. And Tucows and some other people who are in a room down the hall worked very hard to put them into place, and we're really excited to see how the landscape will change because of them. We're looking forward to ICANN issuing breach notices, and I've told them that, and we will see what happens in that arena. But this is what we're really – this is what the community is doing. This is what the contracted parties are doing, and we're looking forward to seeing what data comes out of these

amendments so that then we can figure out what the next steps are going to be.

There will be next steps, as somebody said. The bad guys are always just a half step behind. So, this is one piece that we're using. We're looking forward to what kind of data we get from it because we're going to get some data from it, and I'm really interested to see what that is. And then it'll be interesting to see how we can move forward from that data.

GREG SHATAN: Thanks, Reg. I'll go to Graeme in a second, but both Graeme and you mentioned ICANN compliance, and I'm really sad that I don't have someone from ICANN compliance on this panel because they are the great variable at this point. They have the armory or the arsenal there, and the question, do they have the teeth or do they have the guts, whatever it might take to go to a different paradigm of compliance and enforcement than we've seen. They had baby steps over the years but always said that they didn't have what they needed, and now they do. So, we're all going to be watching with bated breath how things go in that regard. Graeme.

GRAEME BUNTON: Thank you, Greg. This is Graeme for the record. A couple thoughts on that, on a number of things that have been said. First, re-compliance, I think from the contracted party side of things, we feel like we have given them the tools that they needed. I think between the inputs they are getting right now from individual reports that are coming into

compliance, plus the insights that they're going to get from either work like our Compass Project, which provides insights into Abuse in the ecosystem, or from their – the upcoming Metrical project from ICANN, that's going to give them the information I think they need to go take action. And in this case, my sense is the entire community, registries, registrars, AILAC, GAC, everybody is interested in seeing them use these powers.

And so, I think for them that's going to be a lot of impetus, and I am optimistic that they're going to use it. We could spend a long time talking about the difference between the mitigation of Abuse, access to registrant data, accuracy of registrant data, and how those things are combined. I think a key bit for me when I'm looking at this, and the Institute very deliberately stays away from issues of access to registrant data, gets into the world of privacy, and it complicates our work immensely, and so we just put it aside. When registrars are mitigating Abuse, they're trying to do a determination of fact. Is there, say, phishing occurring on this website or in this email using a domain name? That determination of fact doesn't require registrant data. And so, I think, important to keep that in mind.

A number of registries and providers to or services to registries have been very clear publicly that they do all of that PII-free, and so, which is not to say there's maybe no purpose for that data, but there's a lot of work we can do before we need to get into the weeds on that and it complicates our lives a lot. One last piece I'll offer, because it's very interesting to me as the Institute looks at its work and how can we make the biggest difference around DNS Abuse. I put efforts into two broad buckets. There is improving reactive processes. Can we improve

reporting processes so that registrars get domain or Abuse reports that are better quality faster? That's why we created Net Beacon. There's lots of work we can do about improving those reactive processes, and that's really what the amendments are about.

We talked a little bit here, and Laureen was expressing some optimism about sort of, and Ram was talking about some beta work, alpha work, whatever that is, early work on machine learning detection of malicious names. I did a bit of a deep dive on that topic, and I'll share a high-level view of this here. I think I did a presentation at the Vietnam DNS Symposium about this, if anybody's really curious. But briefly, the people who have really published work on machine learning identification of malicious names are mostly from the ccTLD community, and their efforts have not been wildly successful. My estimation is that they were not ready for commercial primetime.

Mostly, I think that is because registries tend not to have sufficient data to build a model that's going to be accurate enough. And that's just the nature of where registries exist in the ecosystem. I think registrars are in a much better place to have that sort of technology to try and identify domain names before they go into the zone and potentially cause harm, which is ultimately what we would like. We want to prevent harm from happening in the first place.

So, then we get into another interesting question, and I would love to talk about this more with anybody who cares, because I care about this stuff far too much, which is, does it make sense to invest in that model, which is going to be a very expensive, complicated engineering undertaking? Or do you just go and look at the services already

provided by the payment processors that retail registrars, admittedly, are already using, and already have very sophisticated anti-fraud detection measures that can be leveraged for relatively inexpensive? And I think that's probably a more practical, more immediate, less complicated technology integration than trying to build our own fancy models. And so, I think there's a lot of thinking that needs to be done before we all sprint down that road. I'll stop there. Thanks, Greg.

GREG SHATAN: Thank you, Graeme. And that reminds me of something that another registrar said to me a long time ago, which is that we really don't care about who is in the sense that we don't use it to keep track of our customers. We have our own databases and our own methodologies, and a lot of that ties in with the payment card industry and their methodology.

And so that makes sense. And I think a lot of the separation between the who and the what, as I said before, DNS Abuse is a what, and you don't necessarily need the who to take care of the what. Obviously, you want to look for recidivists and all those sorts of bad things, but the focus, first and foremost, really needs to be on the what. And then secondarily, we figure out various ways to deal with the who. Steve?

STEVE CROCKER: Thank you. Interesting discussions here. So, a couple of comments. One is that when I think of and use the term registration data, I don't have in mind that it is the exact match of what's collected through what's accessible through RDRS, for example. In my mind, it includes

payment information. It includes the account holder and a variety of other attributes, all of which are, in fact, collected during the right time of registration.

A little word about account holder. Anybody, and almost everybody in this room, I'm sure, has registered a domain name at one time or another. And there's a sharp distinction between the process in which you engage in getting the domain name. You have to have an account with the registrar. You log in. You have a credit card information, usually and then you go and select a name that you would like. And if that name's available, then you register it and you assign who the registrant is, and thereafter, the registrant has the formal authority over the name, but the person who has the effective physical control over it is still the person who has the login account and can make changes in a matter of seconds.

Account holders, remarkably, are not mentioned anywhere at any time. They are sort of a relationship between the registrar and the registrar's customer and don't show up in thin or thick, WHOIS issues at the registry level. So, as I said, take a somewhat broader view. And given that broader view, some of the kinds of questions that one would want to ask for going after patterns become more relevant and important. Again, with respect to patterns, another element is it's one thing to look at the data within a particular registrar.

It's quite another thing to be able to look across the entire ecosystem. And I imagine that, that statement is causing some shudders, like, oh my goodness, what have we done with respect to privacy? What have we done with respect to the investment that it takes to go build this?

And I can say, yep, I understand those reactions, but it is not simply that this creates a gigantic mess. We know how to build those systems. We know how to control those systems. And we could do that if we wanted to. I very much like the comment about making use of the financial industry's controls and mechanisms and processes that they have.

So, you go and you vet a credit card before you accept it, for example, and they have fairly sophisticated tools for deciding whether or not a credit card is good for the purpose at the moment and so forth. So, a number of other things, but those are sort of quick reactions to some of the things that have been said here.

GREG SHATAN: Thank you, Steve. I'll go back to Reg, who has her hand up.

REG LEVY: Thanks. My mind boggles at the number of brute forces attacks our customers would see if their account information was suddenly publicized or even started to be distributed as part of a registration data request. But I wanted to come back to some of the other questions that we were asked. Whether or not ICANN has a role to play in content-based Abuse. And I'm honestly shocked by this question coming from the RALOs.

The picket fence is a thing. ICANN doesn't deal in content, full stop. We can talk about how terrible CCM is and how much registrars and registries are doing to combat it, and we are doing lots every single day of the week, but we don't talk about that here because it doesn't matter

to ICANN's remit. It's completely out there. It's outside the scope, and it's really concerning to see so many calls for ICANN to extend its scope in that manner. I think that the question about whether or not ICANN should publicize its role with regard to DNS Abuse, sure, I think that's a good thing. I think that a lot of people who have heard of ICANN don't know what it is that we do, and it would be great in that sense to have that publicized, but also a lot of people don't know what ICANN is, so if someone were to see tomorrow a news report about how much ICANN is doing to combat DNS Abuse.

I'm not sure that they would even know what to do with that, so I think internally, contracted parties have a job to do about publicizing what we do to combat DNS Abuse, what we are doing, what we are planning to do, and continuing our outreach to groups just like this one to have those conversations with the community, but I'm not certain that people outside of this community really know that much about the fact that they are apparently represented by the people in this room when it comes to their use of the internet.

GREG SHATAN: Thank you, Reg. Jonathan Zuck, speaking of the people in this room, please go ahead.

JONATHAN ZUCK: I'm in the room. Thank you. A lot of what I was going to say that Graeme handled, I think there's a lot of interesting things out there, but the approach of biggest bang for the buck I think is definitely a really strong way to look at this, and I have for the past few years, sort of hung my hat

on the idea of these predictive analytics and things and not really known whether that was real. I mean, there were some papers that came out of .eu, et cetera, about it, and I remember it's because we were all on lockdown while we were having those conversations, and half the conversation was happening in chat, and it was hard to actually know how real any of that was, and the credit card thing, I think, is interesting.

The other thing that pops into my mind has to do with bulk registrations and the correlation between high-volume registrations and Abuse, and I don't know the extent to which that's been studied, and hopefully that's part of this DNS Abuse study that I always forget the acronym for that's going on currently, so I'm hoping if that comes back with something concrete that there's then a way to address that. The question I had, we talk about these contract amendments going into effect in April, but correct me if I'm wrong, the expectation, though, is that people have 18 months or something from then to be in compliance. Is there some number like that that's built into this process? There isn't an expectation that all these bad actors have had enough time to do anything by April. Do I have that right or wrong?

GREG SHATAN: Reg has an answer.

REG LEVY: Absolutely not. They go into effect on April 5th, and I fully anticipate that I'm going to start getting compliance complaints on that day.

GREG SHATAN: So, we have a remote question in the question pod, so I'll ask our remote question pod reader to read the question from the pod.

MICHELLE DESMYTER: Thank you, Greg. We have a question from Olivier Crépin-Leblond. Question to Reg. Isn't a domain name by its very nature in that it can have a meaning or pass on a message already content?

REG LEVY: Maybe, not as far as I'm concerned, but it's a domain name, and the DNS is the domain name system.

GREG SHATAN: I guess even if the domain name has content or at least has meaning, that doesn't mean that all content is suddenly fair game. As much as we would like content Abuse to have a home, obviously, as Reg says, we have remit issues. That could be a very interesting discussion or maybe a very brief discussion depending on how you look at it. But that wasn't today's discussion, thankfully, about content. We have more than enough interesting things that we've discussed.

Just to kind of wrap up here, I think we've learned a lot about kind of domain Abuse in the wild. We've learned about the new DNS Amendments and what they hopefully will bring, both in terms of complaints and compliance and enforcement, the issues and potential solutions that could be found in some ways around registration data,

the particular concerns about time to mitigation, and the fact that the longer a bad act takes place, the more problematic it can be. And I will, last but not least, important that Laureen reminded us that at the end of it all, the ultimate goal is to take an end user and typically to fleece them or mess with them or steal from them in some fashion, or to whether that end user is an individual consumer or it's a business that's undergoing a ransomware attack that somehow got started through a phishing and a successful phishing attempt or spear phishing attempt. And there are real world consequences to this of course, and it's often the end users who in many ways have the most to lose.

Obviously, our ecosystem and our industry, has a lot on the line as well, and hopefully we have turned some corners in terms of the ways that we can help to deal with this. In the end, a lot of this comes down to the fundamental nature of the internet or what I like to call the intermediation disintermediation paradox, which is that you don't know who you're dealing with, but they have the way to get to your deepest and darkest secrets and vice versa. So, you have both the ability to be very close and very far away all at the same time in a way that no other method of communication has ever allowed. It's incredibly powerful, but like other incredibly powerful things that can be used for good, evil, or anything in between. I don't know if anybody else from the panel has any last words. And Ram, you've been so quiet. Anything from you?

RAM MOHAN: I think there've been lots of really good points made. I don't have something extra to add, and that's my general mode, Greg. If I don't

have something to say, then I don't find myself in the position of having to say what has already been said.

GREG SHATAN: I'll learn that from you, Ram, in the future. Seriously, though, I think we've had some really good discussion here. A lot of interesting points. I'm really glad that we had this panel. I'd like to thank each of the panelists for participating in what was a freewheeling and somewhat underdefined panel, and yet I think was very successful in covering a lot of different topics and getting a lot on the table. So, I want to thank the audience as well, here and online, and say that this meeting is now adjourned. Thank you.

[END OF TRANSCRIPTION]