ICANN80 | PF – GNSO: CPH DNS Abuse Community Working Session
Thursday, June 13, 2024 – 10:45 to 12:15 KGL

SUE SCHULER: Welcome to the meeting of the CPH DNS Abuse Cross-Community Workshop. My name is Sue, and I'm a participation manager for this session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior. If you would like to speak during this 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. Please state your name for the record and speak at a reasonable pace. I will now hand the floor over to Dennis Tan.

DENNIS TAN: Hello, Sue. This is Dennis Tan from Verisign, a co-leader of the Registry DNS Abuse Working Group, along with my colleague, Chris Disspain from ID, and my colleagues, co-chairs of the Registrar Stakeholder Group, REG LEVY from Tucows and Luc Seufer from Namespace. Reg and I will be co-emceeing this session today. This is a workshop, and let me just go through some introduction, rules of engagement, if you will, and then we'll go into the substance of the meeting.
So with this workshop, the purpose of this meeting is to build upon the outreach that we've done earlier in the year, the CPH done earlier in the year to community groups such as the BC. So I recognize Mason Cole, IPC and NCSG. During our outreach, we identified some action items of common interest. So we're going to hear from two of those. We'll get

that into a minute. We anticipate that this workshop is going to be one of many in the future. So we want to test the format, so keeping collaboratively constructive tone that will help us go through the session, learn from there, and apply learnings in the future.

The scope of our discussions is going to be on DNS Abuse as the definition of the five broad categories, pharming, phishing, malware, botnets, command and control, and spam when used as delivery mechanism. We understand, we acknowledge there are other online harms. Those are also addressed by some registries, registrars, but that's not the conversation that we're going to be doing today.

As far as rules of engagement, Chatham House Rules, we want to foster, encourage an open, candid conversation, but we don't want to be, so to speak, name and shame, okay? Again, keep it constructive, collaborative, and tone, and so, we're going to have a successful session. The format, we have the agenda in front of us, so we have a two-part program today. The first part is going to focus on abuse reporting. So we have Graeme Bunton from the NetBeacon Institute, and he's going to walk us through the attributes of a phishing abuse report, and how if you want to report phishing to a registrar or registry, what is the exact components and attributes that you need to be looking at.

Followed by that, we have Leticia Castillo-Sojo from ICANN organization from the compliance team, and she will help us understand, describe when you raise an abuse complaint to compliance, again, what are the attributes, the components that you need to consider in order for compliance to consider that as a valid well-formed report.

In part two, we're going to have Farzaneh Badii from the NCSG, and we are going to talk about human rights, and specifically an introduction to the human rights impact assessment, and then, ponder the question whether we can use human rights impact assessment in order to consider that and look at potential future work to build on the DNS Abuse Amendments implementation. So she's going to spend some time raising the education for everybody understanding, and then post questions on potential next steps.

So with that, I'm going to hand it over to Reg to introduce part one of our session.

REG LEVY: Thank you, Dennis. This is Reg Levy from Tucows, and Co-Chair for the Registrar Stakeholder Group of the DNS Abuse subgroup, and I would like to introduce and turn it over to Graeme Bunton, Executive Director of the NetBeacon Institute, to walk us through what good reporting looks like.

GRAEME BUNTON: Thank you, Reg. Hello, everyone. Good morning. This is Graeme. I think collectively between Leticia and myself, we've got about 40 minutes, and so, I'm going to do a chunk on abuse reporting, and then Leticia is going to do a bit on reporting into ICANN, and then we're going to hold some time for questions. I think we're going to try and keep about 10 minutes for that. So store up some good questions. I've got two leftover [Swag] t-shirts, both in XL size, that if someone asks a good question, they can have one. I'm sorry for the lack of selection. Next slide, please. Actionable abuse reports. So the NetBeacon Institute runs a project called NetBeacon Reporter. It is a conduit for abuse reports, and we see

thousands of abuse reports go through this service every month. And from that we can learn some things about what an actionable abuse report looks like, and I'm going to try and share some of that for you today. Hopefully, much of this is sort of common sense and obvious, but for some of you this might be new, and there's some pieces in here that I think we're still working on as a community that might be kind of interesting for discussion. Next slide, please.

So the first key bit to know as you're producing an abuse report is that there's this sort of Goldilocks zone for information that registrars are looking for. They want enough information to make a decision about mitigation, and no more than that. And so, often we see abuse reports that have either too little or too much information, and so, a report will come in saying it's bad, and sometimes there's no domain name or there's no full URL for where the harm is present, and there's just not enough information in that report to do something.

On the other side of that, you often get abuse reports that have like 30-page PDFs often detailing the rich and varied history of perhaps a trademark, and at the very bottom of this thing is the little bit about abuse, and no one wants to have to go through that. So we want to make sure that we're presenting this information front and center. Next slide, please.

And so, this is another important piece, and it's practical, but I'm sure not everyone loves it. But resources at registrars who are dealing with abuse reports are unfortunately finite. And in many cases, an unactionable abuse report can be worse than no report at all because it consumes these finite resources. And if a registrar is working through an abuse report that's incomplete, insufficiently evidenced, then it

delays them getting towards actionable abuse reports. So to the extent that we as a community, can optimize our abuse reports, we can help abuse desks get to more abuse and reduce more abuse faster. And I think that's a really positive outcome we want to remember. Next slide, please.

Linked in this presentation are two pieces of information that I think are useful for the community. One is the CPH abuse reporting guide. This has been out for a couple years now, but it looks at the sort of spectrum of abuse and what information you would need to report into a registrar.

And then, the second is a draft abuse reporting format that we've been working on at the Institute, but also inside of the contracted party house that really lays out sort of the fields and formats that an abuse report should have. And we'll go through some of that at the end of this presentation, assuming we have time. But that's open for a comment. If people have thoughts, please take a look at that and give us some feedback as we're working on that. Next slide, please.

Right. This is an important note on reporting abuse that's only found in an email and there is no visible resolvable website. And really here we're talking mostly about phishing, but it could apply to other harms like the distribution of malware as well. But we see a lot of friction in this process right now. We see it being very difficult to report abuse. And I don't see any easy path for this to get better. And I would love for this community to begin thinking about this and maybe there's ideas that we haven't come up with. And so, if you have some, it would be great to hear about them.

Registrars usually require email headers and the content and email headers are not just the box with addresses in them. Email headers are underlying code in an email that detail really where that email came from. And that's important because email is very easily spoofable. And so, registrars need that information to figure out really where that email came from. Getting email headers keeps getting harder and harder. You can't do it through the Gmail web client, for example. I don't know about the Outlook client. Maybe I'll try when someone else is talking.

And even in the web clients, like if you're in your browser or you're in a desktop application, it becomes harder and harder to find the option to get email headers. You can't just forward the email. You need to be able to forward the email as an attachment, which is not available in every client as well. And so, this is a thing that we need to sort of collectively think about because I don't have an easy answer about how to get non-technical people to be able to discover that information and share it easily. So really we're going to talk more in the rest of these slides about reporting phishing where there's a resolvable website. And that, from what we see going through NetBeacon Reporter, is the majority of abuse reports. Next slide, please.

So there's a couple key elements when reporting abuse. First, there's a domain involved. If there's no domain involved, please don't send it to a registrar. There's another key piece, which is that there is some form of deception of identity, and I'm going to talk a little bit more about what that means in a couple slides. We need an attempt to capture sensitive information. That could be login credentials. It could be some form of other sensitive information, social security numbers, dates of birth, things like that. And then, we need some form of clarity on who or

what is being impersonated so that the registrar can make some choice about this phish. And so, let's look at a couple examples. And so, we can then look at reports for those examples. Next slide, please.

So this is SMS phish that came to me personally, actually, that I subsequently reported to the registrar. And so, you can see that it came in for Royal BK alerts, and there was some impersonation in that text message that it was pretending to be RBC, which is a large Canadian bank for those who don't know. And if you resolve the website, it went to a fake login. So here we had impersonation in the domain name. They're pretending to be Royal Bank. They're pretending to be the Royal Bank in the content of the SMS and the website, and they were clearly attempting to capture credentials. So we have these pieces of information, but we need a couple more to build an effective abuse report. Next slide, please.

So here's what a relatively simple but actionable abuse report would look like for that example. So we've got the domain. We've got the full URL. Note that they've been defanged with brackets, so that that domain shouldn't resolve although looking at this slide, it looks like that one actually might be linked, so I'll need to fix that. And then a short description. And this is what I'm trying to get at here with that Goldilocks piece of information. So it says that I received a text message impersonating a bank, was trying to steal credentials. So there we have the attempt at gathering sensitive information. I've provided here the real website for that bank, which is rbc.com.

Perhaps this domain was registered with not a Canadian or North American even registrar, and they may have no idea who Royal Bank is. And this is a really important piece of information, because not all

brands are globally recognizable. Not all the entities that are being impersonated are known to every registrar. And without including who the definitive source of that is, it's really hard for a registrar to make that choice, and I'll lean into that a little bit more in the next example as well. I've also included here where this harm was visible from. So I could see it on my phone, I could see it from Toronto, and I could see it at this date and time.

We see more and more that bad actors are geofencing or excluding user agents from being able to see the abuse, and that makes it tricky for a registrar. So we want to provide them with the information that maybe they can use to try and see it themselves if they can't right away. And then, I've attached the two screenshots to that abuse report. Next one. Thank you. So this one is a little bit harder, and that's why I've included it. So it was for a swap-blur.TLD, relatively rudimentary login, unclear to me what blur is. I don't know what blur is. So it's an unknown brand to me. There was some form of deception in the content and an attempt to capture what appeared to be crypto credentials. And so, this is some form of crypto phishing, which, boy, we see an awful lot of going through NetBeacon Reporter right now. It seems to be a big problem because there's lots of money in crypto, and there are so many entities in that space that are legitimate that no one has—I'm not a big crypto guy, and so, I've never heard of it.

And this makes it really difficult for registrars to determine if it is actually phishing or not because sometimes these abuse reports are weaponized against competitors. We often see or we have seen cybercriminals reporting their competitors for other cybercriminal activities. And so, this world gets kind of murky. Next slide, please.

So what would an example look like for that? So we've got the domain, we've got the URL. In both the examples I used today, the phish was at the root of the domain name. And so, here we include that other information that we need. Again, we've got the website is impersonating blur.foundation and attempting to steal credentials. That blur.foundation is really important because that is the authoritative entity that is being impersonated. So then we can go and look at that and see what are the attributes of that, what are the colors, what's the logo, what's the branding? Is this person trying to be that. And then, we also include that other relevant information about where it was visible from and when it was visible.

There's also a piece here that can be very useful, which is in this abuse report that was sent in, and this is a real report from a couple days ago via NetBeacon Reporter, that they swore on penalty of perjury that they are authorized to represent the brand and to the best of their knowledge the information is true. What's useful there is that the information related to who is being impersonated, the registrar gets a little bit more comfort that that's authoritative rather than just someone else reporting it or without including that little piece of information because it can be hard to decide with some of these especially in the crypto phishing world which one is real. So those are the two examples I was going to show, I'm going to go through a couple more pieces of information for discussion and then I'll pass it over to Letitia. Next slide please.

So we've talked a bit about, or I've talked a bit about deception here and deception of identity, and this is an idea that I was sort of noodling on recently and we talked a little bit about at the contracted party

summit in Paris, which is, we need a way of disambiguating some of the harms, all of which have some relationship to phishing or some relationship to deception and it's hard to figure out what's going to be appropriate for a registrar or not. And what we're really trying to do here with this deception of identity is really lean into that, this website is pretending to be an entity it isn't, which is different from this website is not doing what they say they're doing.

That could be shipping a counterfeit good, could be not shipping anything at all, but that function, that deception of function is not visible to a registrar. They can't assess whether that's true or not, but the identity piece, they generally can. That information is usually observable to a registrar. It could be obvious to them that it's not the Royal Bank of Canada. Maybe they have to look at their account information or registrant information to be like, right, it seems unlikely that this brand is registered with us. And so, if you're looking at abuse and you're trying to disambiguate it, is this a phish or is it not, this might be a useful concept for you. Next slide, please.

So let's put all of this into a robust abuse report, and this is coming from that draft abuse report in format document. We've got sort of four pieces of categories of information that we're looking for in abuse reports. And this is not definitive, the community hasn't agreed on this, but I think generally, if these reports were coming in this fashion, I think everybody would be happy. So we need the domain name. That is required. You'd be surprised at how often that is not included.

Days since registration is often very useful for registrars, primarily because it helps them understand the risks of suspending a website. And so, if it's only been two days, you can have some comfort that

you're not about to disrupt a huge operation. I'm not going to go through every field, I just want to highlight the ones that are important. Next slide, please.

So the issue info is sort of like the core of an abuse report, and there's a couple pieces here that are really required for every registrar. We need to know the issue type. We need to know what you think it is, whether it's malware or phishing, botnets or spam. I suppose it could be pharming, but I've never seen one. We need the full URL if it's distinct from the domain name, and we want that short description of the issue. And hopefully, those two examples are illustrative of just like sort of how much information is required. It's not an essay, but it is a little bit of work there.

The date and time of observation is important, because in case it takes a little bit to get to the abuse report, or it wasn't filed for some time. The replication requirements are that like the user agent, the geography, those pieces of information are very useful for a registrar if they're going to try and verify that that harm was present. And then targeted entity is really who is being impersonated, and that's that registrar needs to understand who that is. And if they can't, if they don't know that brand, it makes it really difficult to make a choice. Next slide, please.

And then reporter info should be included, who you are and how to contact you in case there's more questions. Anonymous abuse reports get a little bit harder, especially in the context where sometimes they're weaponized. And I think that is just about everything I wanted to go through today. Next slide.

Yeah, there we go. So hopefully, that's helpful in understanding what the components of an actionable abuse report for phishing look like. I

didn't touch on malware or botnets. They're harder. They're more complicated. We didn't have a lot of time today, but we'll maybe leave that for another workshop. So I think with that, I pass it over back to Reg or to Leticia. One of the two.

REG LEVY: Yeah, my job is just to introduce Leticia Castillo-Sojo, I believe. Thank you.

LETICIA CASTILLO-SOJO: Thanks, Reg. Hi, everyone. My name is Leticia Castillo. I am a Senior Director with ICANN Contractual Compliance, which is the function within the org that enforces all ICANN's policies and agreements. Next slide, please.

We do so mainly through the processing of complaints and also through our audit program. We have web forms that anyone in the world can use to submit complaints to us, report compliance issues to us. And these forms also allow us to capture all the data we need, all the information we need to process these complaints and to report on them. Well over 80 percent of the complaints that we receive across all types do not result in compliance cases with the registrars or registries.

For example, last year, most unactionable complaints we received lacked any evidence or even an indication of a noncompliance, referred to country code top-level domains, which are outside of our enforcement authority, or where duplicate complaints. Next slide, please.

On April 5th 2024, this year, we started enforcing the new DNS abuse requirements. And again, most complaints that we received since then were unactionable, largely for the same reasons that I was explaining

before. And something to point out is that we do not simply close these complaints. We provide each complainant with any information that can be helpful, often follow up with them multiple times to try to get the information that we need, but ultimately are not actionable.

And something else to point out is that we do get actionable complaints and we act on them during the same period. The complaints that we receive regarding DNS abuse resulted in over 2,500 malicious domain names being suspended and over 300 phishing websites being disabled, just as an example. So I linked there a blog that we published last week, actually, with this type of information regarding the outcomes of our investigations resulting from actionable complaints. Next slide, please.

So what is an actionable complaint? Our web forms are broken down by complaint type. And each one contains a brief explanation of the type of obligation that we are enforcing and what we need from the complainant. A well-formed complaint includes the complete list of generic top-level domains that are relevant to the complaint. Our system actually automatically closes complaints when they refer to ccTLDs only and sends out to the complainant an explanation and the link to the ccTLD manager contact details. But those are closed. And we also received complaints that list one gTLD and then mention there are many more, but they don't include them all. We need all the domain names to ensure that we address the complaint in its entirety. Next slide, please.

A well-formed complaint also identifies the contractual violation. And we do have explanations on each web form with checkbox. We don't expect everyone to know all the obligations unless they want to join our

team. And then here you have a couple of screenshots of the DNS abuse complaint form for registrars. For registrars it's similar just with questions that are tailored to their requirements. Complete well-formed complaint identifies the type of DNS abuse and also elaborates on how the domain name is being used for the DNS abuse.

We received complaints that say this domain name is being used to conduct a phishing attack. How is it being used to conduct a phishing attack? Is it impersonating an entity, which entity? Displaying a login page, sending emails trying to trick people into revealing data? How is it being conducted? This is information that we need. Next slide, please.

Another important point. We need complete copies of all the correspondence between our complainant and the registrar or registry. And this means that those communications show dates. We are assessing the promptness part of the obligation. So we need to see dates. But they also show the complete communication. We have received complaints that include, for example, a screenshot of an abuse report submitted to a registrar. And then we can see at the bottom a little bit of the response from the registrar. So we will have to follow up with the complainant to get a complete copy of all the communications that took place.

And when it comes to evidence, a well-formed complaint clearly links the domain name to the DNS abuse. For example, if the evidence provide is a screenshot of a login page, the screenshot will ideally show a complete URL. If it's a copy of an email, ideally, we get the email as an attachment, headers even better. But we have received, for example, screenshots of part of an email that does not really provide sufficient evidence. Next slide, please.

And a complete complaint are quicker to act on when the report that was submitted to the registrar or registry was as complete as possible. This goes back to the examples that Graeme was giving. We have seen copies of reports to registrars that say these people are criminals, do something. Sometimes they include a domain names, sometimes they don't. We are assessing compliance with the obligation to take mitigation action when there is actionable evidence. That type of report does not sustain an allegation that the registrar or registry was provided with actionable evidence that the domain name was being used for DNS abuse.

It is also important to give the registrar or registry time to investigate an act before submitting a complaint to ICANN Contractual Compliance. We have received cases, and all these examples are real cases, I'm not making them up. We have received cases where the report to the registrar was submitted within a few hours a complaint was submitted to ICANN Contractual Compliance. When we reviewed the complaint a couple of days later, the domain name is suspended. The registrar just needed some time to be able to assess the evidence and enact on it.

Our web forms also allow the option to select an anonymous submission. This means that we do get the contact details of the complainant, but we will not share them with the contacted party. This is because we want everyone to feel comfortable coming to us and reporting compliance issues. We have instances in the past where they were reporting certain crimes and they were telling us I'm scared that if my details are provided to someone else, the criminals can go after me, for example.

So we want everyone to feel comfortable coming to us, but this anonymous submission do not happen often at all. And for more standard cases, for lack of a better word, we will normally follow up with the reporter, the complainant, and tell them we may not be able to address the complaint completely if we cannot, for example, share some evidence with the registrar or registry that may identify you, and confirm with them whether that is okay.

And another point that doesn't happen often, but has happened, in the advisory that we published, that ICANN org published along with the new DNS abuse requirement, it explained how registrars and registries represent a portion of the DNS ecosystem, and depending on the DNS abuse, one or the other may not be the most appropriate actor to assess, verify, and stop the abusive activity. We have some cases where we have to provide this clarification to the complainant, engage in conversations with them, and this normally delays action a little bit because it takes a few days of back and forth. So this is just a few examples of what we have observed. Next slide.

I think that's all for me. Yes. Thank you. We're ready for questions.

REG LEVY: Thank you, Leticia. We are opening the floor to questions. I do see and will read a question in the chat. Why do I have to go to the police and say, I was at north 41 degrees, 3 minutes, 2.974 seconds longitude west, 73 degrees, 56 minutes, 45.77 seconds, at zero 11 hours, 55 minutes when someone from my 4:00 approached me with a nine mm Browning with a wooden grip and a red Swiss army knife and took my wallet from my right jacket pocket. Why can't I simply say, I was robbed in the city and you investigate?

GRAEME BUNTON: Hi, Reg. This is Graeme. I'm happy to take that one. So I think that question is really trying to get to the issue of providing that replicability information and abuse report seems perhaps like a burden. And so, for there, I'm really talking about what platform you were browsing to that or you were able to resolve that harm on and time of day. And registrars, feel free to chime in if there's other pieces in that replicability that you would also look for. But there are some unknown numbers of jurisdictions that people could be looking at something from and a very large number of platforms. That could be on mobile, it could be an Apple phone. When the harm is targeted, specific device type, none of that is visible to a registrar.

They don't have a technical means of interrogating that information or really figuring it out. And so, if you are unable to provide that information and that harm is geofenced in some way or some other form of blocking, the registrar cannot find it. It's not that they're too lazy to investigate. It just might not be possible. And so, there is a bit of a burden there. Not every harm is geofenced or blocked in that way, but it's useful information. Hope that helps.

REG LEVY: Thank you so much, Graeme. Sarah in the room has her hand up.

SARAH WYLD: Thank you. This is Sarah Wyld from Tucows. Just a quick question for Leticia regarding the forms that you showed where someone can report abuse to ICANN. The check boxes had different terms for different types of DNS abuse. Those terms are well familiar to us, may not be so familiar to others. Are they defined publicly? Thank you.

LETICIA CASTILLO-SOJO: Thanks, Sarah. This is Leticia Castillo with compliance. Thanks for your question. So there is a link in the forum to the advisory where they are defined but we also have that free textbook where we ask them to explain the issue itself. We have had situations—both. We have situations where they select phishing, but the actual description and the evidence proved that it was not phishing. And we explain why. And we have had situations where they selected fraudulent activities. They didn't select phishing, but the actual description and the evidence show that that was phishing, and that's how we addressed it.

So even though we cannot provide all definitions to everyone, because it would be a very, very long form with [inaudible], sorry about that, but we do have that text that is mandatory because we want them to be able to express what they want to express, and we will be able to review. And if we have to follow up and ask questions, or clarification, we'll do so.

REG LEVY: Thank you, Leticia. There's another question in the chat from Raymond. And I apologize, [inaudible] and I didn't say who the question was from before. Raymond's question is, are there any punishments for complainants who do such complaints intentionally for the dubious purposes? Example, a business competitor made such claims against me just to frustrate me. However, the claim is not true. Are there penalties to deter such false claims?

GRAEME BUNTON: Leticia, go ahead.

LETICIA CASTILLO-SOJO: No, go ahead.

GRAEME BUNTON: This is Graeme again. So that's a tricky one. So for NetBeacon Reporter, the service that we operate that allows people to report abuse to registrars where we see problematic reporters abusing the service, so they're reporting repeatedly inappropriate harms to registrars, or maybe it's not an inappropriate harm, or it's just not a real abuse report. They're trying to weaponize the system. We have the ability to effectively ban them from using the service, and we're sort of generally happy to do that, although have not had to do that yet. We usually reach out to the reporter and say, if you don't stop doing this, you won't be able to report abuse via this service anymore.

At the registrar level, I don't know what the answer is, and compliance might have a better thought, but it would seem apparent to me that there is going to be a category of reporter that would be vexatious. They're sending, and I'm sure every registrar in the room has examples of this, where people are CC-ing every person at the company, they're sending 10 emails a day for the same thing, and it's not helpful. It's not abuse. And so, I think it would be perfectly reasonable to have some standard of what you consider to be a vexatious reporter, include that in your terms and service for your abuse reporting function, and then just not accept abuse reports from that user anymore.

LETICIA CASTILLO-SOJO: This is Leticia Castillo with compliance. This does not happen often at all, but we have had situations with what we call abusive reporters. This, for example, entails a situation where they submit multiple reports. I'm talking about hundreds a day, with repeat complaints that

purely being explained in detail how they were incomplete, and using abusive language, either towards us or towards the registrar. When I'm talking about abusive language, I'm talking about insults and really bad stuff.

So what we do is we do send warnings. We follow up with complainants a lot. We explain what the complaint is about. We explain the contractual obligation. We explain what the registrar or registry is required to do, and we also remind them of using proper language during the communications. And after a couple of warnings, if they continue doing the same, we will prevent them from being able to submit complaints. Like I said, that doesn't happen often at all. I can remember a couple of cases, two, three cases, and talking about years, but if we get to the point, yes, we do act.

REG LEVY: Thank you both. Mike Palage, I see you have your hand raised. Go ahead.

MICHAEL PALAGE: Thank you. Mike Palage, for the record. Graeme, a question to you. When you were talking about the difference between identity and function, if you can elaborate on that, and let me give you a hypothetical to maybe walk through. So if someone was impersonating a bank trying to steal their credentials, that's bad and appears to be actionable. What about if someone was impersonating a pharmacy and trying to ship them fake pharmaceuticals or someone was trying to—another person was impersonating Rolex and trying to ship them a fake Rolex.

In that identity and form and function that you were discussing, do you see potential different outcomes in connection with those abuse complaints, or would they all be treated equally or different? That would be helpful if you could expand.

GRAEME BUNTON: This is Graeme. Thanks for the question, Mike. So it's hard to talk hypotheticals without seeing them, which is one of the reasons we include those examples. I think a piece here is useful that it could be both. You could be looking at this website and that it is both impersonating, say, a pharmacy, an actual real-life pharmacy that exists in Canada that might be the brand Rexall, say. And so, they're a fake Rexall, and it's not shipping real drugs or it's not shipping anything. And so, you have both of those things, and when you have that deception of identity, I think you've got enough there to hang your head on for a phish, and you can submit that and hopefully see some action at the registrar, whatever they deemed appropriate for that harm.

My sense would be, and I'm not a registrar, and so feel free to chip in here, guys, would be that if it is a legitimate looking registrar, or sorry, if it's a legitimate looking pharmacy that is sort of not observably impersonating another person and they're shipping either counterfeit or not shipping product at all, that not shipping is not observable by a registrar. Maybe they have a relationship with a reporter that covers that sort of thing as a registrar to make that choice but my sense is that's too opaque, and they wouldn't have enough information to make a choice. Does that help?

MICHAEL PALAGE: Yes, it does. Thank you.

REG LEVY: Thank you. There's another question in the chat from Farzaneh to Graeme. What do you mean by email header? Is it the title? Why is it not possible to download on Gmail mobile? Because of encryption? I hope forward to [inaudible] is actually implemented in many mail servers. And also, the mail server is not DNS, so the email in effect is outside the scope of DNS abuse, I think.

GRAEME BUNTON: Thanks. I'll see if I can tackle that. Boy, I should have included a slide that shows you actually what an email header is and how to get it. And so, if we do something like this again, that's an issue worth diving into. I think the easiest way to answer that is just go search for how to find email headers in a client that you are using. It's just not possible to do in the Android Gmail client on my phone. And so, you just can't get them. You need to be on a different platform to do it. To the extent that a domain is engaged in phishing via email, and so it's maybe the sending domain. Maybe it's the name that is linked inside that email that then resolves to something else. I think that would qualify for what we're talking about here.

REG LEVY: Thanks, Graeme. Chris, go ahead.

CHRIS DISSPAIN: Thanks, Reg. Chris Disspain. I have a question. So [inaudible] in the chat recast his question about the—in essence, it's about the complications or the complications that he sees in filling in forms. This is a question for you, although I think it applies to—it may well apply to all registrars. Have you tested your complaints form with normal people? Have you

actually sent it out on a focus group and said, do you understand this? Can you fill it in? Thanks.

REG LEVY: Just briefly before that, I'm cutting the queue after Lori.

LETICIA CASTILLO-SOJO: This is Leticia Castillo. Thanks, Chris. So I'm not sure what normal people are. I was shown screenshots of our web forms, and I think there was a common—can't remember from [inaudible] related to this as well in the chat. Each box contains a request thus contain also a little explanation and sometimes even examples of what it is that we mean. We understand that this is highly technical and we try to make it as simple as possible while making sure that it's clear what we need. So that's how we try to address it. We try to add explanations. And when that doesn't happen, we will follow up with our complainants and then explain in plain English.

REG LEVY: Thank you, Leticia. Werner Staub?

WERNER STAUB: Werner Staub from CORE Association. I have a question and an explanation for that question. The question is, is anybody working on the complaint relay interface for professionals? Now here's the explanation for my question. All the stuff that I've been seeing, virtually all of it, is coming through interfaces where actually there is a complaint interface somehow by the intermediary, the online intermediary who delivered the content. It may be WhatsApp, it may be a messages app, it may be YouTube, it may be Facebook and so on. In each one of those cases, a domain name has been used as an outgoing link by the

perpetrator. There is a complaint interface, usually hidden. There's issues with that. But let's say if we wanted the complaint to reach us, the obvious way to do this would be to say we offer a complaint relay interface with standardized fields. We have a technical guide with how this can be sent so YouTube can then actually relay a complaint to the respective registrar, whichever system we have, based on their smart interface with the user, so that the user does not need to go through a complex system that we offer and will actually be unable to provide proof of the link because, for instance, if it comes in YouTube, the actual full URL is not even visible on the smartphone where it shows. Maybe they can do a screenshot and they cannot actually show that this comes from a certain ad that was placed on YouTube.

REG LEVY: Thank you very much. There is a question from Lori Schulman in the chat. Regarding reporting forms, it seems like the issue of form design is coming up in these discussions and in our DRS discussions. Would it make sense to adopt guidelines that forms be written in layman terms to encourage a range of users to exercise their access to remedies? Not entirely sure who that question is for, but I agree that probably all of us in the room speak at a different level than most people who receive your standard smish. So it's probably a good idea to give the forms to a family member and say, what does this mean to you and maybe have a UX team. Thank you.

DENNIS TAN: All right. Thank you for that discussion, very useful information. Hopefully, we have the recording, so this is going to be a resource for the future as well. I will make sure to advertise it accordingly. Now we

are going to move to part two, where we have Farzaneh that's going to walk us through an introduction to human rights impact assessment and pose some questions to consider for us on its application for potential work, additional work on DNS abuse practices. Farzaneh, are you with us?

FARZANEH BADII: I certainly am.

DENNIS TAN: All right. Take us away.

FARZANEH BADII: Hi, everybody. My name is Farzaneh Badii, and I'm a member of Non-Commercial Stakeholder Group. For those who don't know the Non-Commercial Stakeholder Group, it is a stakeholder group as a part of the GNSO and we care about human rights, internet governance, accountable and transparent processes at ICANN and on internet governance in general. So today, I'm going to present the idea of doing a human rights impact assessment on the DNS abuse mitigation amendments and get some feedback from you. Oh, wow. These slides are really cool and professional.

So one of the things that I wanted to tell you, so this is like the HRIA that we keep talking about, is a collective work of several digital rights and human rights organizations that some of them are members of NCSG and it is not a new concept. But we just wanted to present it to you and get feedback from the community if we can actually use this tool. So NCSG has two constituencies and advances human rights values, civil liberties, access to knowledge. We want to keep the internet open, secure and global.

So this is another thing that we care about. We also want to bring transparency and accountability to internet governance processes and it's just to emphasize human rights impact assessment is just one tool in the toolbox that we can use to hold internet governance processes accountable and to respect human rights. Next slide please.
So what is this HRIA? So human rights impact assessment is a process through which businesses, governments and organizations assess the impact of their actions and decisions on people and their rights. I can hear a little bit of echo. I don't know if you can hear me okay.

REG LEVY: You're coming through fine. There's no echo on our end.

FARZANEH BADII: Perfect, so it's just me suffering here. Let me just suffer through this. All right. So there is this organization called Business and Social Responsibility that provides a definition of HRIA which is human rights assessment identifies and prioritizes actual and potential adverse human rights impacts and makes recommendations for appropriate action to address those events. It's fair to say BSR is not the only organization that does human rights impact assessment or a range of organizations that do that such as Article 19 and CDT and digital rights organizations in general. Human rights impact assessment assesses the risk to people instead of risks to companies and organizations. Next, please.

DENNIS TAN: Farzaneh, sorry to interrupt, are you still hearing the echo?

FARZANEH BADII: Yeah, we have still remote participants.

DENNIS TAN: Those in the room, can you check your mics are muted. Just take a few seconds there so that people following us on Zoom have a good audio. All right. Let's try one more time. Go, Farze.

FARZANEH BADII: Perfect. So what is this HRIA process? So in order to do a human rights impact assessment, the organization or those who are undertaking the human rights impact assessment of the decisions of an organization talk to the affected communities. So, they hold these consultations with the affected communities that could be impacted by the decisions of the organization and then there is a stakeholder engagement process because they talk—people who undertake these assessments, they talk to stakeholders that have an interest in that decision in whichever field it is and the stakeholders can range from the government to the technical community and civil society members. And then during the HRIA, we identify the human rights that are potentially at risk and then provide recommendations for remedies and mitigation strategies. Great, and the echo got sorted, so perfect. Next slide, please.

Can I do this myself. Okay, great. So, let's look at how this HRIA can work in DNS abuse. So we can look at—so first, we have to look at the affected communities. Actually, this is in no particular order. They can be just done in parallel. So we look at the affected communities and identify the affected communities and consult with them. So in this case, one of the affected communities for DNS abuse contractual amendment might be the registrants, but also might be the internet users. Then holding stakeholder engagement. So engage with the ICANN. So we have the stakeholders at ICANN that we can hold these engagement sessions,

such as the commercial stakeholder group, the non-commercial stakeholder group, the government, but as well as the technical community.

But stakeholder engagement can also identify and move beyond the community that is at ICANN if the people who are undertaking this HRIA see that there are other stakeholders that are not engaged but have a role in this kind of DNS abuse landscape. Then we identify the human rights that are potentially at risk. What are these human rights that could be at risk due to the mitigation mechanisms that are used? For example, it could be privacy. It could be freedom of expression and security. So this is the process that we have to go through and identify the human rights that could be.

And this is like what we're talking about, it's very important to note this, that what we're talking about is the risk to human rights. It might not have happened, but what is the level of risk to the human rights of people when we are implementing the contractual amendment? And then, the last step is providing recommendations for remedies and mitigation strategies. Next slide, please.

So how can we improve the DNS abuse mitigation contractual amendment? And how can we improve the implementation processes through an HRIA? And these are just questions that we want to discuss with the community and we just want to ask, so can using the human rights lens for DNS abuse mitigation processes be successful in terms of increasing the quality and fairness of the mitigation processes? So can we have, for example, provide remedies for the domain name registrants if something happens and lower the risk of making

mistakes? And can HRIA help with clarifying mitigation actions that CPH should take?

So for example, if there is ambiguity for the contracted parties about the mitigation actions or there is ambiguity for the community, can human rights impact assessment help with clarifying that? And then, since the definition of appropriate and reasonable depends on CPH's practices, can this HRIA process help with providing direction on best practices? What is appropriate? What is reasonable? And these best practices, it can be very consultative and consider community member's feedback as well. So it's not industry best practices per se, but it's more of a multistakeholder best practice that we're talking about.

For example, another thing that we have in the amendments is reasonably prompt action. Can an HRIA help the CPH on best practices for defining and implementing terms such as taking reasonable prompt action? And I think that was the last slide. So what are we suggesting? So we are suggesting a set of potential next steps to find out whether HRIA is a good mechanism as an acceptable mechanism to answer those questions and help with implementing the contractual amendments. So one of the things that we can do is scenario-based human rights tabletop exercises. This is what we do in the human rights and digital rights communities, especially when we want to mitigate issues and abuse during crisis.

So we do these tabletop exercises and we put people in different roles and discuss what are the potential human rights that can be impacted, so that we identify the human rights and we also, in the process of that,

we can identify the actors that are involved and the other stakeholders. And we can also reach out to the affected communities or potentially affected communities to receive their feedback on using an HRIA and how we can use it and if we should use it and engage with different stakeholders. And if you like this idea and if we decide that we are going to continue discussing this, we can have sessions with various stakeholders to go through the human rights impact assessment and contextualize it a little bit more.

And these slides and the idea comes from collectively from the Non-Commercial Stakeholder Group. As well as human rights, we care about internet governance, as I mentioned. And also, we can also discuss how we can improve the implementation of the contractual amendments through other tools as well. And that was about it. Thank you.

DENNIS TAN: Thank you, Farzaneh. Now we'll open to questions and answers. So Farzaneh is available and also discussions.

FARZANEH BADII: Okay. So Reg, are you going to read the questions for me or should I just—

REG LEVY: Yes, absolutely. But first we have Edmon Chung in the room.

EDMON CHUNG: Thank you. Edmon, here. Thank you, Farzaneh, for bringing this up. Actually, adding to what you just said, at the ICANN TWNIC Engagement Forum earlier, one of the discussions about this is to potentially add transparency reporting, as well because there's a human rights impact, but also, what has been take down, who has been asking for suspensions, is it coming from governments, for example? And some lines of thinking about transparency reporting might be useful on top of, I guess, some of the potential steps that you mentioned.

DENNIS TAN: Thank you, Edmon. We have Jeff, next.

JEFFREY NEUMAN: Thanks, and thanks for the presentation, Farzaneh. As part of the analysis, I didn't see it in the slides, and I know you were just doing it at a high level, but is the part of the analysis on the human rights impact of not implementing the amendments, I guess, or not taking action where registrars and registries should be taking action, because there are human rights at play if we allowed certain activity to continue? Thanks.

FARZANEH BADII: Yeah. So thank you for the question, Jeff. So as you mentioned, we want to use this HRIA in the context of implementation of these amendments. But of course, one of the things that we also wanted to do HRIA a few years back when we were discussing what are the drastic consequences of DNS abuse. And we have actually looked at some of the—so NCSG

wrote a report on what is the impact of DNS abuse mitigation on human rights. But of course, if the community wants to, we can also look at what can be the risk if they do not implement the contractual amendments.

DENNIS TAN: Thank you. We have Mark Datysgeld next on the queue. Go ahead, Mark.

MARK DATYSGELD: Thank you. Mark Datysgeld speaking. Yeah, I'm following up with what Jeff was saying because we are fighting here malware, botnets, phishing and pharming. Those seem like reasonably clear-cut cases of crime, and the likelihood of somebody practicing one of these tactics, having their human rights harmed versus of the people they hurt en masse around the world every day at scale seems comparatively minor. So if this assessment is to be made, it needs to be made from the perspective of the victims.

That's what we are trying to fight for. We're not trying to fight for the operators of pharming operations. We're trying to fight for the rights of the victims that they affect and for the integrity of their livelihoods, their family, their businesses, their operations, their NGOs. So I don't know, from a certain perspective, to me, this seems to be going about it the opposite way of what we should be doing. It's who are we trying to protect here.

FARZANEH BADII: Yeah, thank you. So actually, the beauty of HRIA is that it can identify those who are impacted and more holistically I'm trying not to leave anybody behind. When we talk about botnets and malware and phishing, but the mitigation mechanisms that we have in place, they might have consequences for the domain name registrant, as well as for the victim. So the HRIA, what it does, it maps those who are impacted. So when we argue that, okay, so there are so many victims, but also, we can identify that there are victims of these malware and phishing, but also, we can identify who else is impacted.

So I think that we can also, through an HRIA, we can provide evidence when we talk about, okay, so these are the registrants that are impacted and these are the victims that are impacted and this is how proportional our mechanisms are without making blanket statements that, okay, so we think that the proportionality is correct or is balanced. It's just another tool to give to the organizations to have a more fair and balanced approach.

DENNIS TAN: Thank you for that answer, Farzaneh. And I think in the chat room also, I think they echo the balance. That's a key term, right? It's a balancing act. It's a risk assessment as to do [inaudible] domain name, what are the implications. And doing an HRIA exercise will help us as a community understand all aspects of one action, what happens, and who is affected, and then from there, make sure those potential victims are less impacted, if that makes sense. So I'm looking at the queue, no hands. I don't see any hands, and no questions in the chat room. So one

final call for questions. No? Okay, so I think we'll wrap part two and Reg will take us through the summary.

REG LEVY: Thank you. And I'd definitely like to extend gratitude to Graeme and to Leticia for walking us through how to best provide actionable reports to contracted parties. One of my key takeaways was to make sure that my UX team goes through our reporting forms for usability, and I hope that reporters also took away some good pointers and some tips. The slides obviously will be available later if you didn't take scrupulous notes. I'd also like to thank Farzaneh for reminding us of the impact of taking domains offline.

One of my primary takeaways was that registrars shouldn't be scared by threatening cease and desist letter, but rather take the time to consider the rights of our customers as well, and to note that everything is a balancing act, and we have to keep that in mind as well.

DENNIS TAN: So thank you, Reg for that. So I think that's a wrap. So thank you, everybody for joining the session and contributing to this discussion. So we look forward to the next one. Farzaneh, thank you for joining from Zoom room. And offline, we're going to continue our conversations as to what are the next steps. Thank you. Have a good day.

SUE SCHULER: We can end the recording.

[END OF TRANSCRIPTION]