ICANN75 | AGM – GNSO: CPH DNS Abuse Outreach
Sunday, September 18, 2022 – 13:15 to 14:30 KUL

SUE SCHULER: Hello and welcome to the CPH DNS Abuse Outreach session. Please note that this session is being recorded and is governed by the ICANN's expected standards of behavior. During this session, questions or comments submitted in chat will be read aloud if put in the proper form, as noted in the chat. If you would like to ask a question or make a comment verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor. Please state your name and your affiliation for the record and speak clearly at a reasonable pace. Mute your microphone when you are done speaking.

The session includes automated real-time transcription. Please note this transcript is not official or authoritative. To view the real-time transcription, click on the closed caption button in the Zoom toolbar. To ensure transparency of participation in ICANN's multi-stakeholder model, we ask that you sign into Zoom sessions using your full name and affiliation. For example, a first name and last name or surname. You may be removed from this session if you do not sign in using your full name. With that, I hand the floor over to Alan.

BRIAN CIMBOLIC: Actually, I think Reg Levy remotely is going to kick things off for us.

REG LEVY: Thank you so much, Brian, and welcome everyone. Thank you for joining us [inaudible] at the Contracted Parties House DNS Abuse Community Outreach session. May I have the next slide please? Thank you. Brian Cimbolic and I will be co-chairing this. And I'd like to thank ICANN technical staff for having excellent operations and allowing me to help out seamlessly from my home. I hope that everyone is enjoying the call. I'm going to turn it over now to Graeme, who will take us though malicious versus compromised domains updates.

GRAEME BUNTON: Thank you, Reg. Hello, everyone. I'm Graeme Bunton from the DNS Abuse Institute. At ICANN73, earlier this year in March, we had what I thought was a really enjoyable plenary on the topic of malicious registrations versus compromised websites. And we got the community to a really interesting place, which was boy, those are different types of harms, and maybe we should be treating them differently as we're trying to mitigate them. But it was very clear that there was a lot of continued work to do on the

topic. And consequently, the Contracted Parties House has spun up a group to work on this issue.

And what we've decided to do is work on a paper for the community that sort of really discusses this issue in more depth. It is not intended to be a sort of operational document or a best practice. What we're really trying to do is get at the nuances of teasing apart these types of harms and what to do about them. To do this, we invited a number of members from the security community, as well as from the ccNSO and began sort of collecting ideas and working on them and putting all of this together. We had really hoped to have this work done ahead of this meeting. But unfortunately, and almost all my fault, we were unable to get the work done on time. And so, we' really hoping to get it done for November.

So I think we have a pretty good idea of what it is we're putting together. And now, we're really just in the process of writing words and then getting that to these people who have agreed to contribute to really edit and refine and help us make sure that this is a useful, good, insightful product for the community. And I think that's really it from my end of things. I'm happy to take questions if there's anything about the work if anybody has any questions. If not, let's throw it back to Reg, I guess. Thank you.

REG LEVY: Thank you, Graeme, and I will turn it over to Alan for the update on Spec. 11(3)(b) for the registries.

ALAN WOODS: Thank you very much. So, the last update we gave on this will probably seem quite similar on this one. And I think it came down to the fact that there was a very short run-up between ICANN74 and ICANN 75. So, I'll give the update of where we're at at the moment. To briefly recap what the effort is, is ICANN approached us knowing that Spec 11(3)(b) is in the ICANN contract for the registry agreement and that registries must have technical analysis and statistical analysis of our zones under our contract. And there's valuable data in how we are getting from Column A, shall we say, of reports that are received and reviewed to Column B. And that is reports that are actionable, evidenced, and ones that we can report on.

And there was an ask for a voluntary program of which we can provide that data that we have as registries to ICANN in order to give some insight into our day-to-day work when it comes to the management of those DNS abuse instances.

So, we have been working through a document trying to make it as homogenous and as easy for as many registries as possible to become part of this voluntary effort. And, of course, and I think I said this at ICANN74, the devil is in the detail. From a high level, it

seems to be something that we… There's a number of people who were on board from day one in that. But as we get down to the detail and how do we make it in a format that is readable, updatable, and useful, that's where we're having a few conversations.

I'm happy to report that since we last talked, of course, we've invited ICANN themselves to the call or to the meeting just to see where we're at in the report and perhaps give us a few more pointers, directions of what they might like to see as well. So, we are having a dialogue with them as well to make sure that it is useful. Again, something that people can use ultimately to give that view between, like I said, Column A and Column B.

So, a very open-ended statement at the end, as far as the aim is to complete the document shortly and seek broad participation in this. And, again, the detail here is making sure that it is as applicable to as many registries as possible to encourage as many registries to voluntary participate in what is, hopefully, a very worthwhile effort. So, with that, I'm happy to take any questions on that as well. Okay. If not, I'll pass it back to Reg then.

REG LEVY: Thanks, Alan. I appreciate that and the work that you are doing on that. And now, I will turn it to Rowena for an update on measuring DNS abuse from the DNS Abuse Institute Project.

ROWENA SCHOO: Thanks, Reg. So, yeah, I am Rowena from the DNS Abuse Institute. To be clear, this is not a project of the Registry Stakeholder Group or the Registrar Stakeholder Group. It is our project at the DNS Abuse Institute. And I'm going to talk to you about something we've launched very recently, which is a project to measure DNS abuse. So, at the Institute, we have this mission to reduce DNS abuse. And as part of doing that, it's really important that we have a really comprehensive understanding of where DNS abuse exists, whether or not it's being mitigated, and what might help in terms of reducing it.

So, to conduct this project, we partnered with an external, independent, academic. And our brief to him was really to find the best way of measuring DNS abuse for our purposes. And we're coming at that from a perspective of registrars and registries getting a comprehensive understanding of this within the areas that they control and with a focus on mitigation.

So, I'm pleased to say that we have launched our first public report this week. I'll put a link to it in the chat. But a big part of this is us sharing our methodology, or, more accurately, I should say Maciej [inaudible] methodology. Maciej is an academic that works out of the University of Grenoble in France. And as part of this reporting, he's compiled a very comprehensive, detailed description of how he's coming up with the numbers that we get

to. So, I would really encourage anyone who's interested in this project to read that, look through the methodology, and talk to us about it. We want this to be really accurate, reliable, independent, academic. So we tried to be very upfront with how we're getting to these numbers.

Our first phase of reporting is focusing on phishing and malware. I've just included some definitions up there for clarity. And yes, I think I already talked about transparency. That was the big point about having the methodology in there so people should be able to look at this, understand it, and if they wanted to, they could reproduce it.

Could I get the next slide, please? Thank you. And so, a little bit more about what exactly we're measuring and what's in our reports. So, our intention with this is to get an understanding of the prevalence of DNS abuse and in this instance, we're looking at just phishing and malware, meaning how much of it is there and where is it.

But we also want to understand the persistence of that abuse. So, it's one thing to know that there's DNS abuse happening. But we also want to know if it's being mitigated and for that subset of mitigation, how quickly it's being mitigated. So, while compiling this project, we have taken a few strategic decisions. And if you have a look at our report, we've kind of outlined what those are and how they align to our priorities. But one of the things to

highlight is that we have optimized for accuracy and reliability. So, we have a relatively small selection of these that come into this project. And we've done that intentionally because we want this to be really well-evidenced.

In terms of where we're heading with this, we really see this as an opportunity to celebrate and recognize good practice as much as it's an opportunity to shine a light on areas that might need improvement and identify different policies and practices with registries and registrars that seem to have an impact on either reducing abuse or preventing it.

Importantly, we've also included a breakdown of compromise versus malicious registrations. And we've done that for a number of reasons. Predominantly, because it's the mitigation action taking place is likely to be quite different in each case. So, we have a definition of compromise which is essentially a benign domain that is being compromised in at any level. So, the reason for thinking about this differently is because you're often going to have a registrar who may be a victim. And we think that probably the process for mitigating abuse when it's been compromised from a registry or registrar perspective is going to take a little bit longer because it's likely to involve communicating with the hosting provider or other parties along the way.

In terms of sort of phases of this report, we've got our first reporting out now, which is high-level aggregate statistics. Again,

I'll put a link in the chat. There's a PDF report, but then, there's also an interactive chart that you can click on, toggle between different things. We want this to be really engaging and interesting for people to go in and move between and understand and interrogate this data.

We will then also be moving towards a more granular approach in the future, which will drill down further into individual registrars and TLDs. And we're also working very closely with registrars and registries. As we launch this, we're encouraging people to come out and talk to us, see their own data, and give us your thoughts. Help us help you. Let us know how we can improve our understanding collectively. As a community, we think this can be something that's really empowering and can create really interesting conversations to move the dial on this. I'm happy to take any questions.

REG LEVY: I see a question from Peter in the chat. Are the data sets from the DNS AI Intelligence Report available?

ROWENA SCHOO: Thanks, Reg. Sorry, I wasn't looking at the chat.

REG LEVY: That's what I thought.

ROWENA SCHOO: So, the data that goes into—Maciej's methodology involves four different feeds. All of those feeds are either free or relatively low cost. And so, you can go and subscribe to them now. There is a list of the feeds in the methodology as well, which I can grab up. It is APWG, Phishtank, OpenPhish and URLHouse. Those are the four feeds that go into this methodology.

BRIAN CIMBOLIC: Thank you very much, Rowena. There's one more question in the chat, a question on the two metrics. What informed the narrowing down to these two metrics? Presumably, and maybe, editorializing his question is to maybe for form of abuse, malware, and phishing?

ROWENA SCHOO: Yeah, and thanks, Brian. So yes, assuming that is the question about malware and phishing. The reason goes back to having evidence that these particular types of abuse have evidence that [Corelabs] could take screenshots and information for and yeah, evidence, essentially.

BRIAN CIMBOLIC: Great. Thank you very much. Are there any other questions for Rowena before we hand things back over to Reg? If not, then, Reg, back over to you.

REG LEVY: Thank you. And I am going to reintroduce our abuse contact identifier tool. We had some people who are confused about the prior names. So, we are now calling it the ACID tool, which just sounds really cool. This is the tool that the Registrar Stakeholder Group has put together following our series of outreach meetings that brainstormed ways to tackle abuse online. It came to our attention that a lot of people don't understand precisely where to address certain types of abuse complaints. And so, often registrars, or sometimes even registries, will receive abuse reports that they can't actually action.

So, this is a tool that will allow people to put the domain name into a box on a web page and get hosting information, email hosting information, if that is relevant, and registrar and registry information as well. Next slide, please. Sorry, previous slides. I'm looking at the wrong slides. I'm looking at the slides that I see.

So, in addition to displaying the hosting provider information, this also says that this is who can best help you with regarding to phishing, malware, botnet, and content issues. Registrars, of

course, can also take action with regard to DNS abuse, but oftentimes, the hosting provider is the best place to start.

So, we're hoping that people will start to use this. We had some complaints that people could not access the prior URL because it had the word abuse in it. So, this is another reason that we switched to the ACID tool. Next slide, please. Based on our preliminary metrics, we have seen an uptick in people who are using this. And I hope that they are not all just me testing it. And I hope that everybody here socializes this as a tool for people who are looking to help us mitigate DNS abuse where they can go and find out who they can best address their complaints to. And I will take questions about this tool if anyone has it.

BRIAN CIMBOLIC: Okay. I'm not seeing any questions. We do have a question, dating back, and actually, Alan may have just answered in the chat. But, Alan, do you want to just take it away? A question from Brian King.

ALAN WOODS: Okay, sure. Thank you very much. So, for those who can't see it, Brian did ask whether or not the plan was, was ICANN or some other going party going to publish the abuse statistic once the normalized format is finished. This is for the Spec 11(3)(b). And the answer is we haven't really decided the formal format. But the point of this is to have it available. It is data to be looked at, to be

used, and one assumes that it would somehow be published on ICANN. I think that is the current hope and thought on that one. So, yes.

BRIAN CIMBOLIC: Thank you, Alan. Thank you, Brian. And with that, we can go back to Reg who can maybe tee things up for our community outreach questions. Reg, if we can hand things back over to you, and if we can go to the next slide.

REG LEVY: Absolutely. So, these are our questions for this particular session to help guide the discussion. What initiatives are your representative stakeholder groups and ACs engaging in outside of just the Contracted Parties House? Do you have any contacts with hosting providers, email providers, content delivery networks? Is there any way that we can help you reach out to these parties to guide those discussions, or hep those conversations?

Are there any areas of concern that you each have? What efforts can we help you with to engage in and investigate and address them? And looking at our current efforts, is there any additional clarity on what next steps might be necessary?

I'll open the floor. I know that there has been some chatter in the chat. But I think Brian is on top of that. And if anyone has any input or questions, I'm happy to take them.

BRIAN CIMBOLIC Don't be shy. These questions really are just meant to help kick things off. I would also add something else to this is that we have in the past had really great success in doing some collaborative projects, be it with the IPC, be it with PSWG. If there are particular areas or particular projects that you're interested in or your SO/AC or SG is interested in potentially collaborating with the registries and registrars, the CPH abuse groups, we're all ears.

We want these sessions to be engaging both to address kind of questions but also to help figure out where we should focus our efforts next. So, that hopefully when we get to ICANN in Cancun, we're in a position to share some additional work beyond just that which we previewed here today. So, all that to say is to invite comments, ideas, brainstorm. And I do see that there is someone that has raised their hand. Werner Staub, the floor is yours.

WERNER STAUB: I wonder if there's any work on their way to work with browser and app developers to make those pieces of software more usable for the end user to actually at least report suspicion about the domain name. Much of the dangerous domain names are

delivered through apps. And most of them are actually messaging apps. In the messaging apps, of course, they are able to reach the most vulnerable people. And typically, they reach the people who are totally unable to report anything, not just because they don't understand what's happening, but also because, actually, the tool that they have available isn't really optimized to do any form filling and that kind of stuff.

But if we enlisted the collaboration of organizations like Facebook with WhatsApp or Telegram or Apple messaging app, it would be possible for people to at least report suspicion. And then we would actually possibly gain access to a large number of cases where mostly the abuse will never ever get reported because it's been highly targeted, and it will not be seen by any pro.

Worse than that, usually, the type of abuse that is now taking place is such that it only delivers the malicious payload to the intended target, to the party that has a certain profile in terms of http headers, in terms of what kind of IP numbers that they're using, what kind of device that they're using. And if a probe tries to figure out what is happening there, it's just not going to get anything just innocent responses or an error response or anything like that.

So, the only way to actually get data that matters is to make it easier for normal users who are not experts to actually report at

least a suspicion and not a long, long list of questions about what proof that the found as to whether this is really a dangerous domain name. Statistically, more could be done to actually then check if this is worthwhile and to investigate further, but probably using machines rather than manual interaction with reporters.

BRIAN CIMBOLIC: Thank you very much, Werner. And I think Graeme Bunton has raised his hand to answer that question.

GRAEME BUNTON: Yes, I'll see if I can take a crack at that. Thanks, Werner, for the input. So, I think you're talking about a really interesting problem that's pretty difficult to solve, which is how do we make it easy for people to report abuse across in an industry when abuse is technical in nature. It's certainly not a CPH endeavor. But, again, it's a project of the DNS Abuse Institute where we launched a tool called Net Beacon in June that attempts to solve this problem to make it easy to report abuse in a way that you don't need to know who the correct party to mitigate it is. It's going to solve that for you. And so, it's available. Anyone can use it.

But the next steps for that, having solved some deliverability issues over the summer so that we can now be pretty confident when we take reports with Net Beacon that they get to where they need that work is to continue to drive usage. And so, part of that

work is reaching out to the email providers and browser vendors because they do have good quality data. And they have an ability to capture that in a way.

And so Net Beacon is built with APIs. And so, I'm reaching out to them at the moment to see if we can build some sort of connectivity to close that circle, to ensure that quality abuse reports can be easily submitted through this service from a variety of sources. So, hopefully, that helps there. I think it would be very difficult to do across each individual contracted party. And so, baking that into a centralized service like Net Beacon makes sense.

There was another question in the chat about the Institute's Measurement Project about including other feeds. And I'll answer that because I'm talking right now. The short answer is yes, I think we're open to that. The slightly longer answer is that to a certain extent, it's up to our vendor or [Core Labs] who are really the ones who have the extremely detailed understanding of the quality of lists that are available. And so, it's up to them to determine whether a new source of data is up to scratch, if it's sufficient to really begin measuring other types of harms or adding an additional feed into the ones that we've got. We're open to it. It's just got to be high quality. Thank you.

BRIAN CIMBOLIC: Thanks for that. And Alan also wanted to chime in on the question from Werner.

ALAN WOODS: Thank you very much, Brian. Another thing I will add to that, and I think that goes.... Thank you very much, Werner, for that question because it goes to that concept of other efforts that we can talk about. And I think harking back to the SAC 115 document that was very clear on the concept of things like interoperability. And I think where a platform is being using as I suppose a delivery mechanism that is not being caught by traditional feeds, I definitely think that is something that we should and would welcome people to talk to us. Yes, we have lost connection. Oh, all right. I'm going to pause. Cheese and coffees in the lobby.

UNIDENTIFIED FEMALE: You know any good jokes, Alan? Any good jokes?

ALAN WOODS: No.

REG LEVY: I have a TCP/IP joke. Does anyone want… Are you ready to hear a TCP/IP joke?

ALAN WOODS: Yes. So you can hear us, right, I'm assuming?

REG LEVY: Yeah, everything's fine for me.

ALAN WOODS: Then, we shall resume. But we will come back to you for the TCP/IP joke. Sorry. So, just in case, I'm just saying that from an interoperability and being able to work together, this is not necessarily something that falls in the hands of one or the other. This is something that we need to work hand-in-hand with the platform providers because it is not… It's difficult for a registry to deal with taking down on platforms, etc. Therefore, we need to work with those platforms. So, if people like Meta, people like Twitter, TikTok, I'm just not gonna do an entire list. If they wish to approach us and say, "Hey, this is something we'd like to work with you on," then we welcome, absolutely welcome that. And it will be a very interesting discussion and, hopefully, a very good project that might come from it.

BRIAN CIMBOLIC: Thank you very much, Alan. We did have another question, and I apologize. I don't know the exact wording because the Zoom, when it rebooted, we lost the chat queue. But it was from Ephraim at Article 19. And I'm sorry if I poorly paraphrase this, but

essentially, the question was whether or not the CPH abuse working groups would consider conducting a human rights impact assessment in designing the tools. And I think it's a great question. And it's not something we've specifically discussed.

PIR, Public Interest Registry, the company I work for, we've conducted a human rights impact assessment. And it was an excellent process, and we learned a lot. So, it's something that, Ephraim, perhaps we can potentially invite you to participate in one of our meetings and have you walk us through that process because I think it's a great idea. And the whole idea behind these sessions, in particular, is collaboration. And I think that would be an excellent use of our time to have you come explain what that entails.

One administrative matter, too, because there is a lot going on in the chat. If you have a question that you specifically want read in, or if you have a question that you want us to specifically address, rather than just making a comment, please designate that question, colon, that you are asking a question that you want us to specifically respond to.

UNIDENTIFIED FEMALE: Yes. Brian, just a quick technical note. I think for the folks in the room, the Zoom room is down. So, we can't see anything that folks who are participating remotely may be putting into the chat.

We do have full audio capabilities in the room though as far as I can tell. So, if you have a question that you've placed into the chat and you want to get connected to read it out, we can make that happen even if you are remote. And folks who are in the room, we're going to old-fashioned raise our hands like Volker is doing right now.

BRIAN CIMBOLIC: If I could just real quick, Sam, maybe we can ask our virtual MC Reg, who presumably has not lost the room. If there are other questions, maybe, Reg, could we ask you to jump in. But in the meantime, Volker.

VOLKER GREIMANN: Yes. Thank you. And I am using those tools already. So, the DNS Abuse Institute's reporting function is very helpful. And the DNS abuse tool, ACID tool, is also very helpful for us in forming reports where to go next. I was just wondering if we could maybe also integrate those two tools. So, for example, if we say that if a registrar would like to see their abuse contacts in the ACID tool reflected as the form for the DNS Abuse Institute's reporting functionality, that could be very interesting just to make sure that the reports that we are getting through the ACID tool are also well-formed and actionable for us. Thank you.

GRAEME BUNTON: Thanks, Volker. This is Graeme. Just to respond to that, that's a great idea. That should be trivial and easy to do.

BRIAN CIMBOLIC: Excellent. Thank you, Volker. Thank you, Graeme. Reg informs me there are currently no questions in the virtual queue. Are there any questions in…? Mason, go ahead Mason Cole.

MASON COLE: Thanks, Brian. Mason Cole with the Business Constituency. I was just wanting to know, does trusted notifier work into any of the plans either on the DNS AI or on contracted party efforts on DNS abuse? Where do things stand on trusted notifier for you guys?

BRIAN CIMBOLIC: So, I can just chime in that we published a document on this. And so, maybe that's a good point of feedback coming into future meetings is referencing back to what we've already put out because it's a great question. And, obviously, trusted notifier is a topic that a lot of people are focused on right now. We did put out a publication called "The Trusted Notifier Framework." And so, we'll put that… To the extent we can access the chat, we'll put that in there. Graeme, did you...?

GRAEME BUNTON: For the DNS AI on this for Net Beacon, we have baked in the ability for registries and registrars to both label and flag reporters, so that they can log in and say, "I trust this person." I hesitate to use the trusted notifier capital T capital N, because I think that comes with quite a bit of weight. But it does have the mechanism that registries and registrars can flag users. And any future abuse reports from them come with those flags or those labels so they can triage those tickets faster. It's sort of our first step at working towards that sort of functionality. Thank you.

BRIAN CIMBOLIC: Thank you, Mason. And I see Ashley had her hand raised, and then Paul McGrady. Ashley.

ASHLEY HEINEMAN: Yes. I just wanted to respond a little bit because I think you raised a really good high level point as well, probably unintentionally. And Brian touched upon it, which is we do a better job of getting out… It's great for us that we're doing all this work. But if it's not anywhere that's accessible, and even cross-referenced, we need to do a better job. And we actually created, at least in the registrar-stakeholder group, a communications subgroup. So, we're hoping to get better with that so you all have easier access to some of this work that we're doing. So, we'll keep you posted on that. Thanks.

BRIAN CIMBOLIC: Thank you, Ashley. Paul McGrady.

PAUL MCGRADY: Paul McGrady here. Graeme, I'm sorry to put you on the spot, but a while ago, some of us raised the issue that the phishing report in the Net Beacon, it presupposes that the phishing was occurring by email, and it was asking for emails to be attached. Some of us have clients who have phishing by impersonation. And we report that to the appropriate host or registrar and they're usually the same outfit. And the response we get back within an hour is, "Well, we can't confirm that," even though we've shown them, like, they're using our branding. It's got a field to put in your personal information. It's not us. What do we think that personal information's being used for? Well, it's phishing, right?

And so, at one point, you had indicated you guys were going to look into maybe having a drop-down option for phishing by impersonation. And I was just wondering is there any forward progress on that? And if not, I just want to encourage you to keep it at the top of the stack because I think that would be useful. We're hopeful that we can consolidate the complaints with you. Then, patterns will emerge, and maybe that will put some pressure on folks to take those kinds of complaints seriously. Thanks.

GRAEME BUNTON: Thanks, Paul. This is Graeme. I'll respond to that. So maybe briefly, a Net Beacon update. And then, if people don't mind, I'm conscious it is not a CPH endeavor. And so, obviously, we're part of the CPH in some fashion. But it's not a CPH project. So, if anyone wants to get mad at me for that, please do. We launched Net Beacon in June. We spent much of the summer making sure, as I was saying, that if we're taking an abuse report, we can get it to where it needs to go with confidence. And I think we've achieved that.

We now have a long list of bugs and improvements that we're working on cleaning up. And one of them, one of the first ones, is a sort of new, phishing form that does presume less that it's email, that it could be SMS-related phishing, or IM-related phishing or some form of impersonation that's happening. That goes to the dev. It's sort of ready for them in the next probably week or two. And so, I would expect that to be fixed pretty quickly. Thank you.

BRIAN CIMBOLILC: Thank you, Graeme. We have Lori.

LORI SCHULMAN: Hi. This is Lori Schulman from the International Trademark Association. And I want to thank all of you for really, I think, just speeding up and diving into these voluntary efforts. It is very much appreciated by the intellectual property community for sure. With that being said, I still have some questions about scalability and it also goes to the point that Ashley made about the marketing, so to speak, of these efforts because I know INTA is certainly ready to evangelize for you in terms of using these tools.

And we have 30,000 individual brand protection professionals who I think would love to jump on board. But I do have a little bit of concern about how this would scale and are you ready. Are you ready for our community to dive in that way? If you are, we're here to help in terms of we could potentially do workshops. We could invite you to our meetings. There's a lot we could do, include you in our publications, include links to the tools. So, I think there is a lot. And there's other global associations, as well, or regional associations that could certainly help in the intellectual property field, for sure.

Law enforcement will have its own issues and ways of approaching this. But I'm really talking about the private sector here. But, as I said, I have some concerns about sustainability and stress testing the system because I wouldn't want to go to my

members and say, “Hey, here it is. Use it,” and then it’s not ready for us. So, can you speak a little bit to that?

BRIAN CIMBOLIC: Thank you, Lori. Great question. And Reg has raised her hand that she’ll respond to this.

REG LEVY: Yes, thank you, Lori. This is a helpful question. And I think it sounds like you are primarily speaking about the DNS AI reporting tool and also, potentially, the ACID tool.

LORI SCHULMAN: I am.

REG LEVY: I can’t speak to DNS AI’s scalability and how much they have stress tested. We have not yet done a full stress test of ACID tool. So, that’s probably not yet ready to be slammed by thousands of your members at the same time. That said, please do, if anybody asks, refer them to it. And we have many resources available on both the Registry Stakeholder Group and the Registrar Stakeholder Group websites.

Some of them are cross posted on both. And we could probably do a bit better job of doing that. For example, the trusted notifier

framework that Paul asked about earlier. So, please do socialize those. And if you want us to appear at such events to help socialize them, I am happy to do that. I am sure that many others are happy to do that here, as well. I don't know if anyone can see the chat, but Ashley just tossed into it that we are not a marketing powerhouse. That is correct.

This is our marketing. We are telling you. But we are trusting that if you think there are people who need to know that you will help us get that information to them. If that means sending them a link, great. If that means you sending me an email and saying, "Hey, can you draft something that I can send in my newsletter to my whole community," that's great too. We don't know the avenues that you necessarily have. So, we're happy to work with you on that score. And if there's anything that we can provide, please do let us know.

BRIAN CIMBOLIC: Thank you, Reg. We have Graeme and Owen to also respond.

GRAEME BUNTON: Thank you. So, admittedly, we were pretty quiet on outreach over the summer for Net Beacon because after we launched it, we learned some things about its usage. And we really wanted to make sure that we cleaned up some quirks and are now really turning towards that outreach. So, we're ready. We have stress

tested it. We have done a whole bunch of very robust security testing on it. And I think we're ready for that volume. And so, let's have a conversation about how to do that.

BRIAN CIMBOLIC: Thank you, Graeme. Owen.

OWEN SMIGELSKI: Hi. So, Lori, I'd like to respond not as anybody who is involved with the Net Beacon but as a user of it. And I've got to say please do tell people to submit abuse complaints to there because the reports that come through there, the requirements they have, and the documentation they want are wonderful. It makes it a lot easier for my team to get a complaint in there. Quite often, we'll get complaints that won't have the domain name or a screenshot, or they'll send us a 25-page letter about all their trademark rights and everything. We just want to know the information. So, the Net Beacon gets all that info that we need back quickly on it. So, please do that. I absolutely love seeing those reports come in.

GRAEME BUNTON: Thanks, Owen. That's lovely to hear. I really appreciate that endorsement. I should just add that tool is free. It's not a paid service in any respect. It's open to anybody around the world. Please go check it out and thank you.

BRIAN CIMBOLIC: Thanks for that, Graeme. Lori.

LORI SCHULMAN: Yes. I wanted to follow up with two points, if you don't mind. So, yes, and I will commit after hearing today that you're kind of ready to do this. I'm very happy, and I'll do this as a follow-up to my association. We do have a newsletter, and we have a website. So, we'd be very happy to find streamlined ways of just giving this message that isn't too complicated for our members and happy to cooperate. This is a great effort, and this is what we need to be doing. There's, I think, policy development issues inside ICANN that are still [inaudible] discussing. But these kind of efforts, I think, just launch us in a direction we need to go. And so, I'm just super happy to see all the progress.

And then, a point that you raised, Graeme, about the Trusted Notifier system and about the weight it carries. And that raised some questions for me, understanding that the Trusted Notifier system itself is contractually based and it's a bilateral sort of agreement. And so, using that terminology outside of that bilateral relationship would be problematic. But I like this idea and in catching onto this flag or we're going to flag people. So, maybe you might think about definitely not calling it a trusted

notifier but a frequent filer. And actually, think about how you market that to frequent filers.

Some of the governments, the patent and trademark offices, identify frequent filers. And they know which law offices are high volume, which law offices produce better documentation. And there is kind of a system in place of I wouldn't say trusted. I would just say known or well-known partners. So, to build that in and to start emphasizing that, I think, would even be… it might be even more beneficial than sticking to the trusted notifier model, which, as you said, it brings a lot of weight, legal negotiations, and heavy responsibilities on both sides in my view in terms of if it works well.

BRIAN CIMBOLIC: Thank you very much, Lori. So, I think it's a great point, Lori. And at its core, it's not—I think it's obviously a determination that each operator needs to make for itself. But in my experience, trusted notifier is really just an expression of confidence in the reporter. It's a formalized agreement. That's true. But that's not to say that you can't get referrals from entities that you don't have a trusted notifier relationship with that you still [don't] give some sort of expert designation. So, I think it is important that we don't think of things as either trusted notifier or nothing.

There's room for shades of gray on this. And we, oftentimes, do rely on expert third parties that we don't have a formal trusted notifier relationship in and, ultimately, action the domain name in the same way as if it was part of the trusted notifier. But, again, it's really a determination on our part understanding the expertise and the confidence in the referral. So, I think there's room for both in that system. Alan, yes, please.

ALAN WOODS: Sure. And what I will also say as well just maybe to add a little bit on, what we are looking for is evidence and that we can substantiate the claim made and that we are the appropriate party to receive. We don't need to have an agreement in place. Once we get the evidence, it just sometimes helps grease the wheels to get that evidence to us by having those agreements in place, as well. So, really, it doesn't matter if you're a person on your home computer a large company. If you provide us with the evidence we need, well, then, that is enabling us to be able to be far more responsive.

BRIAN CIMBOLIC: Thank you. We've got Volker I think on this, and then Brian and Mike.

VOLKER GREIMANN: Yes. I think by working through Net Beacon, trusted notifiers can get a certain reliability score as well. Look at large retailers that have certain information on their site. Nine out of 10 registrars see this as a trusted notifier for example. That would be helpful information for other registrars that don't know that notifier as well. You also have to take into account that certain notifiers, they are 100 percent on-point when it comes to certain types of complaints. When reporter A sends me malware complaints, I know that's a take-down. But they suck when it comes to other things.

So, flagging someone as a trusted notifier, maybe we need some more granularity on that as well for certain types of responses. But, ultimately, for us, dealing with the intake of the abuse complaints, having such a tool that aggregates those complaints, that investigates and provides additional feedback beyond what the reporter sends you that upgrades the information, that makes sure that the reporter includes everything that we need to know when we take down, it's so helpful for us when we investigate a complaint because it saves so much time and already gives us the information that we need without having to come back to you and ask for it. So, for us as a report intake provider, registrar, this is so helpful.

BRIAN CIMBOLIC: Thank you so much for that, Volker. Brian.

BRIAN KING: I wanted to commend and thank the DNS Abuse Institute for putting together the abuse reporting tool. I think it's fantastic, and I'm really encouraged to hear Alan and Volker say that it's been really helpful for their registrars. I wonder if you can speak to are all registrars automatically connected to this, or do registrars need to do anything to start receiving feeds or reports from that system, or how does it work? Because I think the IP folks would be helpful in our engagements with registrars to encourage them to do whatever it takes, if Graeme or Rowena can speak to what registrars need to do to get hooked up.

GRAEME BUNTON: Sure. Thank you, Brian. Again, I'm aware we're talking about a particular initiative from a particular party. So, if someone is upset and wants to whisper at me to chill out, or we can take this offline, I'm happy to do that, of course. Briefly, Net Beacon is capable of reporting abuse to all gTLD registrars because, by contract, they're obligated to have a contact. And so, we've worked pretty hard over the summer to ensure that as we're getting stuff, we can get it to where it needs to go. So, registrars don't actually need to do anything to receive reports from it. If they choose to create an account and sort of claim their registrar, this is where they get more power and utility out of the service,

which is where they can select the enrichments where we include other data or prioritize those. They can, then, flag notifiers.

And that's, again, getting towards that sort of reporter relationship piece. That's where they get the value if they've actually… They can set sort of custom endpoints or consume by API. If they want to do that sort of thing but they don't actually have to. We do not currently have the ability to report on ccTLDs or to ccTLD- specific registrars. That's on the roadmap, or is that ecosystem considerably more complex than the Gs? But hopefully, we'll get there in the not-too-distant future.

ALAN WOODS: I'll jump in here, and I'm not going behind the curtain too much. But I was personally from a registry point, because it's very registrar-focused at the moment with the DNS Abuse Institute. But the back end of Net Beacon is [CleanDNS,] is something that Identity Digital actually uses. So, if a report comes through Net Beacon, we're really happy to know that that will be coming into our system and then moving along pretty quickly because of our backend provider.

So, from my point of view as a registry, although, of course, we will always try and defer to the registrar because they are the most appropriate and the first instance, we are getting those reports, as well, so that we can start moving them along faster.

And, again, it's all about that TTL and making sure that we get that down. So, that's something that I'm personally quite happy and excited about with Net Beacon, as well. I'm, obviously, depending on how [CleanDNS] itself grows, that's also something that would be available I would expect.

UNIDENTIFIED MALE: Thank you.

BRIAN CIMBOLIC: Thanks, Brian, Graeme, Alan, and Mike. Mike Palage, did I see? Did you have you hand up? No? Okay. And we don't currently have any… Brian, is that an old hand in the…? Yes. Any other questions, comments, ideas of future collaboration? If not, I think can call this a wrap. Thank you very much, everyone. Thank you for participating. We look forward to talking with you in Cancun.

SUE SCHULER: Thank you. We can end the recording.

[END OF TRANSCRIPTION]