ICANN78 | AGM – ccNSO DNS Abuse Standing Committee Community Update
Wednesday, October 25, 2023 – 1:30 to 3:30 HAM

CLAUDIA RUIZ: Hello, and welcome to the ccNSO DNS Abuse Session on Tools and Measurements. My name is Claudia Ruiz, and I am joined with my colleagues, Bart Boswinkel and Andrea Glandon, and we are the remote participation managers for this session. Please note that this session is being recorded and is governed by the ICANN expected standards of behavior. During this session, questions or comments submitted in chat will only be read aloud if put in the proper form, as I've noted in the chat. I will read questions and comments aloud during the time set by the chair or moderator of this session. Interpretation for this session will include English, Spanish, French, and Arabic. Click on the interpretation icon in Zoom and select the language you will listen to during this session. If you wish to speak, please raise your hand in the Zoom room, and once the session facilitator calls upon your name, kindly unmute your microphone. Before speaking, ensure you have selected the language you will speak from the interpretation menu. Please state your name for the record and the language you will speak. If speaking, be sure to mute all other devices and notifications. Please speak clearly and at a reasonable pace to allow for accurate interpretation. This session includes automated real-time transcription. Please note this transcription 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 in to Zoom sessions using your full name.

For example, a first name and last name or surname. You may be removed from the session if you do not sign in using your full name. Thank you, and with that, I will hand the floor over to Nick Wenban-Smith. Thank you.

NICK WENBAN-SMITH: Thank you for the introduction. For those of you that don't already know me, my name is Nick Wenban-Smith and I represent the .UK ccTLD. But for the purposes of this afternoon, until the community finds a better qualified chair, I'm chairing the DNS Abuse Standing Committee of the ccNSO. So we're always very interested in getting more participation. If what you see this afternoon makes you more interested in these sorts of topics, then you're very welcome to join our group. We have a really nice, diverse set of participation in the DNS Abuse Standing Committee. But there's always room for more. It's an extremely cooperative and collegiate group. So it's one of the more fun parts of some of the slightly dry ICANN topics.

So this afternoon, you can see on the slide the agenda. And I will try to chair it. By the way, I find it quite confusing when it has the timing on the slide in UTC. But I think it's actually a full two-hour session, I believe, even though it says it's not on the slide. So sometimes, a bit like with DNS abuse, you shouldn't necessarily believe what your eyes are telling you. It's a metaphor, shall we say.

So we're starting off just with a little recap of some of the resources that the DNS Abuse Standing Committee are providing. And then we'll go on to the main bulk of the session. It's going to be a deep dive into different

aspects of measurement and tools for this interesting and important topic of DNS abuse.

And then finally, because I think one of my great interests is, obviously, to make, first of all, my ccTLD very safe. Secondly, all of the ccTLD is very safe. But I think we need to appreciate, as a community, that there's ccTLDs and gTLDs. And we have a lot of things in common when it comes to measuring and preventing abuse. And there's no point us doing something in our own cc bubble, silo, not including the gTLD along. So I want to have a bit of a chat, because the registry stakeholder group of the GNSO do have a working group on DNS abuse. So I think we want to bring them into the conversation and to have some observations about things that we can better do together in terms of global cooperation. So if we can move to the next slide.

So just a little recap for newcomers. The ccNSO does not set policy for ccTLDs in this area. We are limited in our remit to setting policy for specific IANA function type activity. But when it comes to policies for individual ccTLDs, we are not here to set policy. You can see, therefore, if you are, as a lawyer, you know, I'm very keen to look at the terms of reference. And I spend all my free time looking at my own terms of reference. But just to summarize them here for you, it's much more about information and insight sharing. We do want to showcase examples of what goes well encountering DNS abuse, and sometimes we learn more from the things that didn't go quite so well. So raising your awareness and understanding and talking about what works. And I myself am very open and freely admit I make mistakes. But it’s much more useful to learn from other people’s mistakes sometimes. So that’s part of it as well.

And I think this open and constructive dialogue has been actually—I take no responsibility for it, but as it happens, I feel that the group of people who work on this have been particularly collaborative, and it's been a brilliant working group. And, yeah, ultimately, although we don't set policy, we would like to give the community the resources to enable them to be better custodians of their own TLDs, and fundamentally, we want to help people make the Internet safer, but specifically where it's DNS-related and to help people mitigate this impact. So I've moved to the next slide. Yeah, still has the same timing. And as you can see, COVID has been kinder to some of us than others, but we'll leave it there when it comes to the shots. So I think, as you see, there's resources, the main tools and measurements, and then the dialogue. Next slide. I will hand over now to David.

DAVID MCAULEY: Thank you, Nick. Hello, everybody. My name is David McCauley. I'm employed by Verisign, and I participate in the ccNSO by virtue of our management of the .cc country code top-level domain, and I very happily participate in the DNS Abuse Standing Committee. And so I'd like to say welcome to my ccNSO colleagues here. And to visitors. What we're going to talk about now, for many of you, will be a reminder of what we're doing on a subgroup of the DASC that I lead, which is known as the repository subgroup. And as Nick said, our subgroup, too, is not about policy. It's not about code of conduct. It's all about the creation of effective collaboration tools for ccNSO, for ccTLD managers to effectively address DNS abuse.

We're first going to talk about one initiative we have underway, which is the creation of a repository of useful information. And for that, I'm going to turn to my colleague, Adam Eisner of CIRA, the .ca country code.

ADAM EISNER: Thanks, David. Hi, everybody. As David said, my name is Adam Eisner, and I work for .ca, the Canadian country code, and have been pretty active in enjoying my time with the repository subgroup so far. I wanted to spend a few minutes talking a bit about the repository and resource library that we've been building within the DASC. Go to the next slide, please.

So, really, the goal is to create a dedicated space that ccTLD managers can look at and use as a resource for information sharing. It's managed by a small content board consisting of some DASC subgroup members, where we come together a couple of times a month to review content submissions. And through my next couple of minutes on this, I'm really going to spend a lot of time encouraging everybody to please visit, please have a look, please submit. I think you're going to hear a few times today in the work that we're doing that this is new, this is just getting started, and we're trying to get as much community participation as we possibly can. And that is certainly no different for this repository of information.

So, the board meets a couple of times a month, to review what is coming in content-wise. And then, you know, our goal is to also have a regular cadence of briefings as well. So, you know, in that sense, being

able to provide information and also encourage submissions. So, next slide, please.

So, in the repository, there's really four general content categories. They're pretty straightforward. Presentations, reports, tools that the community can use, policies, definitions around issues, items related to DNS abuse, and then various articles and commentary. So, when you visit the repository, it will largely be divided into those categories. And in here we've included a link in QR code to the actual repository, so you can visit at any point. Next slide, please.

The qualifications or criteria for what could be submitted is pretty wide, pretty large. Certainly, obviously, being related to ccTLDs and abuse is important, but that said, it can be certainly a wide range of topics within the ecosystem. It can pertain to registrars. It can pertain to registries, themes, threats. It doesn't really matter. We're trying to cast a wider net rather than cast a thinner net, certainly to start as well as we build up a base of resources for the repository.

When it comes to format, that also is pretty flexible. I think, generally speaking, any kind of format that you would typically be consuming your news and information in, we're willing to review. It's very simple, high-level kind of requirements in that sort of way. And all submission types are certainly welcome there. Next slide, please.

Not a great slide to necessarily indicate what it looks like, but here's a grab of a few different things that kind of indicate what is in the repository already. So it certainly ranges from things like, I don't know, interviews with the DNS Abuse Institute to presentations that have been made within the ccNSO in previous ICANN meetings, and then a wide

range of things that run the gamut between definitions, tools, and other things. So it's really a good wide range of different information types that speak to both different audiences on different things, and we're just going to continue to try and keep building a base of content in these categories from there. Next slide, please.

And then finally, in terms of the process, it's pretty straightforward. It's kind of low-tech. So really, we're taking submissions at a specific email address that is included in the slides here. And as much information that can be provided as possible is appreciated when we're taking a look. So you can see a few of the things that we're ideally looking to receive as part of a submission. Just to make it a little more straightforward and easy for us to be able to categorize when we review and if we put up onto the repository. Within a couple of weeks of anything being submitted, we will take a look. The group will all, you know, individually will review and then come to a consensus on whether we should be putting it up as appropriate content for the repository or not. And as I've mentioned a couple of times here, certainly, any and all submissions are welcome. We're really trying to build something new and interesting here for the community. Next slide, please.

And so just, final bit, obviously, it is a brand new thing that was launched at the last ICANN meeting and is just getting started now. And I'm going to pass it back to David to talk a bit more about one of our other initiatives in the mailing and contact list. Thank you.

DAVID MCAULEY: Thank you, Adam. So as you can see, on the timeline that's present on the screen, at ICANN 78 we intended to announce our email list. We just missed the target. We are creating an email list tool modeled on TLD ops which I'm going to speak about just for a few minutes. We have completed our work but we have yet to pass it off. We will pass it off to the DASC at the end of this meeting. The approval process there will be pretty quick and then we will make an announcement to the community about this TLD ops like list in the coming month is my estimate.

And by the way, it indicates that our colleague Fernando Espana of .us is making this presentation. Unfortunately, he has fAlan ill, and I'm happy to pick this up in his place. Next slide. The dedicated e-mail list, let me talk about the broad parameters of this list. As I said, this will be modeled on the TLD Ops list. And so what we're creating here is a collaboration tool. This will be e-mail. It will be meant to be a closed list for ccTLD managers. And, of course, being e-mail, there can be no guarantee of confidentiality as people can forward things, et cetera. So we won't make that kind of promise. But we do expect this to be managed and handled and used pretty much like TLD Ops is, which has become, in my opinion, and I think widely seen this way, a very successful tool for security issues. So this list will have -- and in our announcement in the coming month, we will indicate that the list will have an information-sharing element to it, generally speaking. It will have periodic indications of contact points where people have consented to have their contact point published. And that will allow for offline or off-list discussions where people can see that someone may have experience with an issue that they're first encountering, et cetera.

It's that kind of tool. It will be open to up to four subscriptions from anyone ccTLD. That ccTLD will manage who exactly it's going to be among their group. You can have a role account. In other words, you could have three people named and one role, such as the chief security officer or the chief policy officer, whichever you wish. It's a list that we expect will appeal to legal, technical, operations, security, those kinds of things. But it's up to the TLD manager. And we'll give quarterly contact list summaries. So, again, it's much like the repository itself. It's meant to aid in collaboration amongst ccTLDs in addressing DNS abuse. Next slide, please.

So you can see here that information about how to learn about it. Now, this is a little premature in the sense that we will put this information in our announcement. And so this slide is good. And if you're looking at the slides, please pick up on it. But we will bring this to your attention as we announce this to our ccNSO and ccTLD manager colleagues. Next slide, please.

There is information on how to subscribe. That will be in the announcement also. I don't think you can read that here. It's pretty simple, just signing on to an e-mail list. We will have particular notes about the nature of the list, the intent, et cetera, and how one qualifies for it, et cetera. Next slide, please. This is our small six-member repository team, Fernando Espana, Adam, Mary Uduma, Jordi Iparraguirre, Diego Ernesto Luna, and myself. And we are here to respond to the ccNSO community and ccTLD managers.

And so as we launch these collaboration tools, our request to you is please do use them, as Adam was saying. Send us things you think are

of interest. Sign up and take a look at the mailing list. We can nudge this various ways in order to get it launched and make it responsive appropriately. You know, it's in formation stage in a sense. We will be responsive to requests for changes, you know, to make it work right. That's the whole intent. So with that, could we have the next slide. So, Nick, I will hand it back to you now.

NICK WENBAN-SMITH: Thank you, David, and thanks to all of the repository team for those good work to provide the resources to all of the community. So the next part of this session is the meaty part of it in terms of a deep dive into tools and measurements. I will have a few introductory comments, and then we have four different presentations. And I guess from a sort of personal note, like, we're all trying to measure the same thing. So why is everybody coming up with these different answers, and what's a boy supposed to do with all of these sort of conflicting messages about how much criminality and abuse there is out there, but yet when we try to do things, it seems that it's always someone else's problem. So we're going to have some quite interesting different perspectives. So if we move on to the next slide.

The survey we did in the fourth quarter of 2022 asked a very simple question, and this is self-reporting, and so we're not necessarily following the same definition of what we mean by abuse. There are various sort of caveats around conclusions that you could draw from the survey responses. There are two conclusions which we took, and one of them is not so much of a surprise, which is that the ccTLDs participating in the survey indicated that they've got very low levels of

abuse, which is, I think, nice confirmation of what we already knew. But the other very interesting point here, and it's the bar on the right, is that 38% of the people who responded to the survey put this box, which says they were not sure how much abuse they've got. I think that raised a few eyebrows in the working group, because that troubled us, that if you're a responsible ccTLD manager, you should have a handle on this sort of thing. So that's partly the genesis of this whole session, to give people more insight and information into the tools and managements. It turns out measuring abuse is not a formulaic thing. It requires a degree of sophistication, and there's lots of nuances in the data. But the objective of this session is to – the idea would be to repeat the survey, probably in the fourth quarter of 2024. It would be very interesting to see if we get the same answer, having had these sorts of sessions, because I want to see that people now affirmatively know the low levels of abuse that they've got, rather than not being sure. So if we can move on to the next slide.

So this is just a couple of introductory thoughts from me to get everybody's minds working in the right way. I think managing a safe, secure, trusted ccTLD for your national community is one of the fundamental objectives that a responsible registry operator will have. And here I'll put this quote. Apparently – I looked it up because I know the quote, but I didn't know who made the quote, but apparently this is wrongly attributed. But the quote is still a good one, which is that if you can't measure it, then you can't manage it. So you need to be able, I believe, to have a baseline idea about what levels of abuse you've got, and then you can try different policy implementation, and you can see whether that has its achieved effect. But this is to make this more

evidence-based, and to actually have proper management of something, you need to understand what it is and to track it and manage it. So then, if you move on to the next slide.

So this is my favorite psychological experiment. Hawthorne was this factory, and they got some fancy consultants in, and they wanted to see what the worker productivity result was of different levels of light. So they changed the light bulbs, and they wanted to see if the workers were more efficient. And one of the interesting things they found was that even if they didn't change, so the control people who didn't have any light bulbs changed, everybody, whether or not they'd had their light bulbs changed, all were more productive. And essentially, because people know that they're being observed, that in itself changed their behavior. So I think when we come to tracking and measuring, what I'm thinking is there's an analogy here that actually tracking and awareness and measuring stuff modifies behavior. So with those thoughts, I think we move on to the substantive tools and measurements presentation. So I think we're ready to hand over, I believe, to Rowena.

ROWENA SCHOO: Thank you very much, Nick, and thank you to the DASC and the ccNSO for inviting us here today. It's really nice to speak to you about this important issue. If I could go to the next slide, please. So I'm going to give you a brief introduction as to what the Institute is. Then I plan to talk about measuring DNS abuse in general and how that's typically done. And then I'll tell you a bit more about our project to measure, which is called COMPASS. I'll speak a little bit about the chAlanges and then some final thoughts on where does this leave us. So if we could –

thank you. So the DNS Abuse Institute was created by Public Interest Registry about three years ago. And they have done that in pursuit of their non-profit mission. They are a charity in the US and they need to invest the money they make running the registry into things that are good for the public and good for the internet.

So when we think about DNS abuse, we talk about technical abuse. To us, that means phishing, farming, malware, botnets and spam when used as a delivery mechanism. And the purpose of creating the Institute was really to take this issue that had been bubbling up around the community, find the areas of complexity, pull those into the Institute and try to solve them for everybody across the whole ecosystem. So we don't focus on .org, we focus across all TLDs and indeed all registrars as well. Something that is important that we often forget is to say that everything we do is free. So I'm not asking for money or selling anything to you today. We are presenting a bunch of free resources that you can use to better understand DNS abuse. We also have an advisory council and that has a selection of people from the community on it, including Bruce and Nick as your ccTLD representatives. Next slide, thanks.

So when people measure DNS abuse, they generally follow this kind of process, which is when you're setting out to measure, you're thinking carefully about what your purpose and what your focus is going to be. Most studies on this take an input from what we call reputation block lists. These lists are generally created for the purpose of network blocking rather than studying DNS abuse or anything at the domain level. There are hundreds of them, if not thousands. There's varying levels of quality and generally they don't tend to come with evidence associated to them as well. So the next step naturally tends to be that

you need to do some sort of cleaning, depending on what you're trying to achieve. So at a minimum, if you're interested in unique domains, for example, you're going to need to reduce that list of URLs down to a list of domain names. Sometimes there's IP addresses and things like that on there as well, so you need to focus in on what you want to measure. It's important to note that there is no known list, to my knowledge, that doesn't have some level of false positives on it. So it's important to be thinking about that and accounting for that as you go through a study.

Then depending on what's trying to be achieved, there's varying levels of analysis and then decisions made about the presentation of data. And there's really difficult decisions to be made all the way along this process, and that can result in different outcomes and different numbers. So next slide.

So I wanted to talk a bit about our project, COMPASS, and the decisions that we have made throughout that process. So when we started the Institute, we really needed to have a very in-depth understanding of DNS abuse to inform what we were doing at the Institute. But we also wanted to provide that information to the community so that they were empowered with information about what's happening in their zone. We needed to do that at a level that was granular enough to be informative, and we wanted it to be accurate and comparable over time and across the industry. This feeds into our efforts to write best practice documents, and ultimately our aim is to reduce abuse at the DNS level. We have a few principles that underpin our project, and those are transparency, credibility, independence, accuracy, and reliability. So how we went about that is we contracted with an independent researcher and a lab. So Maciej Korczynski, who is also sitting in this

room, is the researcher that we chose to work with. He's a professor at Grenoble University. He runs something called KOR Labs, which is a university spin-off. And we work with Maciej to achieve the technical analysis that underpins COMPASS.

There's a few important decisions that we made that you should be aware of when you're looking at our data. So firstly, we are not measuring all harm on the internet. We're not trying to measure everything bad. We are optimizing for accuracy and reliability over comprehensiveness. We do this so that it allows us to do some benchmarking over time and across the ecosystem. One of the other decisions we made is that we were really interested in not just how much abuse there is, but whether it was being mitigated. So measuring for mitigation was an important next step for us in this project. We also really wanted to understand the type of registration. So whether the domain name was registered for the purpose of DNS abuse, or whether it was a benign registration that was associated with a level of compromise. So usually that's compromise at the website level. So if someone creates a website, they don't update their content management system, there's a vulnerability over time, and a malicious actor exploits that. All of these are really important for understanding what mitigating actions should be taken, and what the risk of collateral damage is.

This is just to give you a bit of a visual of what I was explaining in terms of us focusing on quite a specific segment. So of all the harm that exists, anyone can only measure what is detected and reported. Sometimes people talk about this as an iceberg principle, where there's an mountain under the water that no one is aware of. It's obviously quite

difficult to measure things we don't know about. From what is detected and reported, we tend to focus on security threats. And then for this project, we also made the decision to focus specifically on phishing and malware, making the decision that that was what KOR Labs felt that we could currently reliably evidence.

So this is quite a busy slide, but bear with me. It's also a visual representation of what is in our methodology, which is available on our website. It was really important to us that people could understand how we were coming up with the numbers that we came up with. So we used this process that I mentioned earlier, and this type of process is really something that a few projects pioneered a few years ago, one of those being DAAR. So we take a selection of inputs for us. That's APWG, PhishTank, OpenPhish, and URLHaus. We go through a process of deduplication and cleaning, and something that's really important to realize is that some of these lists have 70,000 or more URLs with the same domain names. There's reasons for this. Some attacks create new URLs any time a visitor goes to the domain name, whether they're human or a bot. So that's something to be aware of. Another important step is that KOR Labs remove what we understand as special domains. So these are domains where action at the DNS level would be inappropriate. So we remove the special domain providers, URL shorteners, file sharing services like docs.google.com. This is a list that is a lot of manual effort to maintain. It's something that KOR Labs make publicly available, and to the extent that people can help us update this, we would appreciate that.

The next step is the analysis that takes place. So there's a fingerprinting of the site, evidence is collected, screenshots, DNS records, those bits.

Then there's a determination as to whether the registration is malicious or compromised. KOR Labs then revisit the website to check against that initial fingerprinting and all of those attributes that they've taken in at the point of block listing. We're doing this to check whether mitigation has occurred, and that happens at short intervals initially, five minutes, 15, 30, then one hour, two hours, out to six, and then every 12 hours for 30 days.

Finally, we take all of this and we summarize it into our public reporting, and we make a number of decisions around how we present that information. So, for example, our reports come out monthly. We include tables about abuse rates per 100,000 domains under management, and we do that on a monthly basis. We make a few exclusions in that data. For example, if you have really low levels of phishing and malware identified for that month, we'll take you out to guard against the potential for false positives, and we only focus on malicious registrations in our public reporting. Next slide.

So when measuring data, there are a few chAlanges that apply to basically all projects. All measurement is going to be a best efforts attempt. The quality of the input is an issue here. There are issues of false positives. We have a blog coming out shortly that gives a deep dive into one of these chAlanges. Phishing simulations, turns out they look a lot like phishing to the outside world. Sometimes they end up on block lists, and there's gray areas as well that are also an issue.

One of the chAlanges we have in terms of the presentation of the data is that often it's quite skewed. It's rarely a normal distribution, and this can also be difficult when you're dealing with small numbers, which is

the case often for ccTLDs. We have a suspicion that these source lists also have some geographical biases, and finally we have some specific chAlanges with ccTLDs in terms of enumerating their zone size and new registrations, but we're working on some collaborations to help improve the accuracy there as well, which any ccTLD can participate in. Next slide.

So this is just a short example. I've put a lot of information on here if you want to go back to the slides, but one of the decisions that we make in our reporting is that we don't take an average of abuse over 12 months, but we have a process of redaction. So if you have a high level of abuse per 100,000 domains under management and you're in one of our tables, but you're only there for one month, we won't publish your TLD. We'll redact you unless you've been there for four or more, and the reason that I chose to implement this is if you're a small TLD like this example and you're the target of a malicious attack in one month, you mitigate everything, you could still end up with a really, really high abuse per 100,000 domains under management that seems not right and not reflective of what is going on in that TLD. Next slide.

So what is a ccTLD to do with all this information? Think carefully about what purpose you're trying to achieve, interrogate how numbers are reached, and really look into that and what those decisions result in. Be aware that false positives do happen. And finally, help us, help us make this data better. We would like to hear from you. We have free dashboards for every TLD and every registrar. You can come and sit down with us and we can take you through that in great detail and help you understand how we get to those numbers. And just this week we've launched access to that through a login, so we can get you signed up so

you can look at that any time that you want. We're nice, we're friendly, those meetings are really helpful for everyone, we really enjoy them, so please do reach out to us. I'll put details in the chat. Next slide.

And this is just a short example. So this is the dashboard for .org, which Public Interest Registry kindly said that we could show. It's giving you raw numbers of abuse for phishing and malware in one month on the left, and then on the right it's normalized for competitive domains under management. Next slide.

This is an example of compromised and malicious and that breakdown. And then on the right you can see mitigation rates. And, yeah, you can go next slide. I think that's all. Thank you. We're around to ask questions. Come talk to us. And also Marcia gave a presentation to Tech Day. If you didn't catch it, you could watch the recording if you want to go into more detail.

NICK WENBAN-SMITH: Amazing. Thank you very much, Rowena. Store up your questions, we'll keep them. I think it makes sense to go through all of the presentations and then to have the questions and answers together. So that's how I prefer to manage that, unless there's any violent objection. Next up we have the DNS Research Federation. I'm very pleased to welcome Nathan Alan. The floor is yours.

NATHAN ALAN: Thank you. It's nice to be here with you all today. So if we could go to the next slide, please. So today I'm going to cover who we are at the DNS Research Federation, the project that we've been involved in regarding

ccTLDs. That will just be a high-level view, and then we can carry on from there. So next slide, please.

So we are a UK nonprofit organization, and our mission is to increase the understanding of the DNS and the Internet and the impact on cybersecurity policy and technical standards. And we achieve that through a number of different ways, one of which is through assisting education and research into the DNS. We like to help those people that we talk with have access to data, and we've done that through creating our own DNS analytics platform, which I'll talk briefly about today. And lastly, we like to engage in technical standards and try and understand what standards are being proposed and how we can help with that. Next slide, please.

This is our website, and this is showing some of the research projects that we've been involved in recently. The TLD league table is what I'll be briefly touching on today relating to the levels of abuse across different TLD types. We also have research projects relating to Internet standards and the measurement of RPKI adoption. And finally, we also have undertaken recently some work on blockchain domain names. Next slide, please.

So, the report that we recently published earlier this year was for the ICANN BC, and that wanted us to look at specifically abuse rates, European ccTLDs. And what we wanted to understand, we wanted a way to rank each of the European ccTLDs, but also all ccTLDs, gTLDs, and legacy gTLDs. And we approached that at a very basic and high level where we took a number of different feeds available in our analytics platform, and for each report that we saw for a particular domain name,

we counted that domain name once. And that's important in this measurement because we wanted to create an abuse rate that was based on the market share of that TLD. So we wanted to compare apples with apples, if you will, so a single domain name would act as a single—And we would divide that by the registration size of the TLD to give us an abuse rate for that. And as we've heard, our data showed that ccTLDs do have lower levels of abuse, and European ccTLDs came out with the lowest levels of abuse. And so the report that we were asked to produce was around understanding those TLDs that were, I guess, at the top of the list, that were the least abused. What were they doing that meant that there was lower levels of abuse on their TLD compared to others? Next slide, please.

This is our website again where you can feel free to go to our website, go through the research zone to check out the data that we've presented. We're also here at ICANN with a stand, so if you'd like to see a deeper dive into how we were able to produce some of our rates and how we came to some of our conclusions, then we'd be more than happy to share those with you as well. Next slide, please.

So the report--and this will mainly target European ccTLDs just because that was the report that we presented-- obviously there was low rates of abuse against those, and we found that there wasn't really a silver bullet in terms of they were all doing one thing that meant that they had low levels of abuse. They all were very proactive in their different approaches, and the study not only took into account abuse that we had seen in the data that we had collected, but we also took guidance from CENTR, a CENTR study that helped us understand some of the measures that these registries were taking in terms of understanding

their customers, randomly spot-checking registrants and their data, and making sure that the data is accurate and correct. So as I said, there was a number of different measures that were taken. There wasn't one single thing. Next slide, please.

And this is, again, just illustrating where you can find links to that report, and I think that might be the end. Next slide, please. And this is, again, a deep dive into the table that we've produced, and we've separated out the number of reports that we had seen on our platform that related to just the domain name, the host name versus the URL. So there was a breakdown there of those numbers. But as I said, we've got a stand here, so please do feel free if you'd like more information about how we came up with those, or if you would just like to see it yourself, we can give you a demonstration of that. Next slide, please. And those are the links to the various resources that I've mentioned today. And that's it for me, I think. Thank you.

NICK WENBAN-SMITH: Amazing. Thank you very much, Nathan. So already we have two free, I believe, resources for all of the people in this room who want to try and measure and track abuse levels. Slightly different methodology, broadly speaking. So get them both for free. And now we have Samaneh. I just prefer to call her Samaneh. I can't imagine why. Welcome, Samaneh, to take us through some of the ICANN.org work on measuring abuse levels. Again, slightly different lens through which they measure things. But Samaneh, the floor is yours.

SAMANEH TAJALIZADEHKHOOB: Thank you, Nick. Hi, everybody. I'm Samaneh, and I'm here today from ICANN's security, stability, and resiliency team. And this presentation is made by me and my colleague, Dr. Sion Lloyd. Next slide, please.

So we made this presentation slightly different from typical ICANN presentations. We start with a question of what are the researchers in the committee are trying to do when it comes to DNS abuse. When you look at most presentation and work of the researchers, most people are trying to measure DNS abuse by some means of creating metrics. Next slide, please.

So now we want to look into why these metrics matter, and why are we trying to always create better and better metrics to have more preciseness and more reliability. In this slide, you can see one set of data presented in two different forms. The plot on the left shows you DNS abuse raw counts, and the plot on the right shows you percentages. Each dot represents a TLD. Now, what you can see is that you can get a very, very different picture if you just subset your data into different metrics. The point is you can get the same data presented in two different ways and get different conclusions looking at the same subset of data. Next slide, please. That is why it's very important to pay attention to how we create metrics and how we report on metrics. What can make metrics better? Higher quality data, data with less false positives. As Rowena already pointed out, when it comes to certain type of datasets, for instance, the reputation block list data, there is almost no data without false positives. But there could be taken some measures to improve the quality of aggregating the data and cleaning the data. Then there could be more precise definition of what is

measured. For instance, we often hear spam as an unsolicited email versus spam as a delivery mechanism for malware. These are very different two types of data which is often not paid attention to in our community. The third point is incorporating measurement errors, acknowledging the data that we are talking about. Let's say if we are talking about a certain RBL or other lists, what that list can represent and cannot. If we are showing the community a plot, this plot is presenting this, but this plot is incapable of presenting why. Next slide, please.

Who we are within ICANN? We are ICANN's security, stability and resiliency research group. We kind of act as a think tank and work on any research related to security of Internet identifiers. The goal of our group is to increase reliability of the metrics that measure security around the identifiers. We either start from metrics that are already existing and improving those, or we develop metrics from scratch. DNS abuse is one topic that we work on. We work on other metrics related to security. The core of our work is to use data and scientific methods to answer questions that are of interest to our community. Next slide, please.

Next slide, please. This is an overall work schema of our group. Having in mind the main goal to create better metrics to measure security and insecurity around Internet identifiers, the black big box there, we collect data and we create metrics. This is an example for some of the projects that we work on. The data that we are collecting are often either data that represent the domain landscape, which is, for example, zone files that contain information about the domain namespace. It could be also the routing data, which is for IP addresses, or reputation

block list data. We do take measures to make all these datasets an improved version of them. For example, when it comes to the reputation block list data, we have a project that works on creating methods to evaluate the RBL lists that are used by us, our community members, and also members of the panel here, and the scientific community, whoever uses these lists.

When it comes to zone files, we developed methods to basically represent the parts of the zone, domain namespace that are active versus parts that are inactive, such as parked domains, which in the research paper we published this year came out that almost around 20% of the zones each day can be not active domains.

Then we also have projects which work on data that are reported, for instance, unsolicited emails reported to anti-phishing working group, which is different from the anti-phishing working group feed that is used in the community. We worked on classifiers to classify threat types and improve the metrics there. When it comes to measurements of abuse, we have the DAAR project, which is the Domain Abuse Activity Reporting Tool, which most people in this room are familiar with, and we are having projects to measure uptimes of maliciously registered domains reported on the reputation block lists. This is, in a nutshell, what the SSR group does research on. Next slide, please.

Now, I will spend a couple of next slides focusing on the Domain Abuse Activity Reporting Tool specifically, because we are aware that some ccTLDs are participating in this project. We would like to have the feedback, but also present on the progress of the project. Next slide, please.

You may already know this, the DAAR project aggregates data from RBLs at the TLD level and combines it with the count of domains from zone files, similar to what my other colleagues presented here. The project is in operation since 2017. It creates monthly reports, in general monthly reports that are public for TLDs, and TLDs have daily access to their own abuse numbers through the ICANN MoSAPI. Since the project is in operation for the long term, it is also capable of creating longer term time series trends and analysis. Next slide, please.

These are the ccTLDs that volunteered and are participating in this project. Once again, we thank the ccTLDs for coming forward and for all the feedback that we received on a bi-monthly basis. When it comes to ccTLDs, in addition to the daily stats that they receive from the ICANN API, they also receive a monthly report that is individualized, customized to the ccTLDs. Next slide, please.

This is an example of the kind of data that is in the monthly reports. For the sake of this presentation, I anonymized the plots and there is a .cc in the plot that is an exemplary ccTLD. There is no attribution to any of the ccTLDs who are participating in the DAAR project. The plots on the screen show the proportion distribution of the four threat types that we use in the DAAR project over the gTLD space and the ccTLD space. As you can see, it seems like the spam as a delivery mechanism is concentrated more in the gTLD space than the ccTLD space. Same to other threat types. Next slide, please.

This is an example of the visual that is in the tailor-made ccTLD report that is sent to each CC, to each country code technical contact that we have. If you see in the plot, there is a .cc dark blue spot that is a ccTLD.

The plots allow the participants to identify themselves in comparison to the rest of the space. This specific visual shows where each ccTLD is, sorry, where each TLD is in general, how they perform in terms of percentage of abuse per threat type. This is a visual that is used for phishing malware, spam as a delivery mechanism and botnet command and control. The size of the circle indicates the amount of abusive domains that are reported on the RBLs. Next slide, please.

Again, this is another example of the visual that is in the ccTLD reports that ccTLD managers receive. It allows for each ccTLD to compare themselves to their peers. This shows the time series trend of percentage of abuse from October 2022 now. As you can see, there are differences between how ccTLD abuse concentration evolves over time than gTLDs. I want to make a disclaimer that the preciseness and the reliability of our data relies on the most basic indicator that we have, and that is the number of data points that we have. For the gTLD space, the number of data points, i.e. the number of gTLDs that we have in our data are way more than the ccTLD, because our data of the ccTLD is limited to the volunteers who are volunteering to this project. The more volunteers we get, the more precise any statistics or visuals that we create will get. Next slide, please.

This is what the project can and cannot do. The project as it is in the current state cannot report on any level other than TLD level, so we don't have domain level metrics or registrar metrics. The project can create historical analysis and can help operators get a sense of where they stand in comparison to the rest of the space, but because the data is anonymized and aggregated, it cannot directly be pinpointed to a

specific domain that is listed. We do not have public ranks of registries at the moment about their security concentrations. Next slide, please.

Why would ccTLDs be interested in joining in DAAR? Since this is a crowd of ccTLDs and we have very active ccTLDs participating. The first and the most important point is that the reports that we send to ccTLDs will allow them to compare themselves to their peers. They can see where they stand, even though they do not see the precise numbers, but they can get a sense of how they are doing on a monthly basis in comparison to their peers. It will help us to give you better metrics if we have more ccTLDs, more accurate metrics. By participating in the program, you can give us feedback. We modify the report up to now with whatever feedback we've received, so you can actually influence even the visualizations, the methodology, and towards what you think would be better, so you will receive better results. And you can use the metrics we create for your own internal monitoring as a ccTLD. Next slide, please.

What is next is that the DAAR team, which is located in my group and is going to be led by a colleague of mine, Sion Lloyd, is working on developing the next version of the project. The next version, it's been already almost a year that we are working on defining the requirements and specifications of the project. In the next two slides, I'm going to talk about how this project is going to evolve. Next slide, please.

The next version is going to be a measurement platform. It's going to be modular, meaning that every sometime there will be a module released that has some functionality. The first module will be including what the current DAAR is plus registrar metrics. The next module, which will be approximately released in nine months' time or a year, will include

metrics for registries and registrars. It will have a dynamic dashboard, which will allow for users to see domain-level data as well as registrar and registrar-level data. It will allow for searches and visuals. There will be an API connected to it, which will allow for researchers to pull the data if need be for the academic community or whoever would like to get the data from other channels other than the web dashboard. What you see in this slide is rather the longer-term plans that we have for the project. Of course, we would like to have more ccTLDs participating to hear what would be more beneficial for them. In future modules, not module one, but future modules, we would like to also have measure uptime metric, meaning the amount of time that a domain is from the registration time that is up for malicious activity to distinguish between malicious registrations and compromised domains, to add if a domain is parked, to add a tag to that, and so on and so on. To also add a module that gives risk scores for predictive measures.

These are the plans for the future. I again ask for the ccTLDs who are interested or have questions to come forward to give their feedback, or if you are interested to participate, I have a slide that is hidden that has the steps of how a ccTLD can join the DAAR project, and you can always approach me afterwards.

NICK WENBAN-SMITH: Thank you very much, Samaneh. So, another sort of, does it from a slightly different perspective, but also I guess ccTLDs need to opt into the DAAR. So, it would be very interesting to see over time more ccTLDs joining in with the DAAR collecting, and it's just another illustration that there's yet another way to cut up the data from different sources. I've

got a few questions of my own, but let's go on to the final presentation of the metric session. So, this is a slightly different, but also extremely interesting CC specific. So, I'm talking about the collaboration between DNS Belgium, the .BE registry operator, and .NL SIDN. So, Thijs and Ronald, I'll hand over to you for your section now.

THIJS VAN DEN HOUT: Thank you. My name is Thijs van den Hout, I'm a researcher with SIDN Labs, which is a research team for the .NL registry. And with me is Ronald Geens, he is the Chief External Relations at DNS Belgium for .BE, and we will present the collaboration between the two of our registries on trying to detect malicious domain name registrations. In this presentation, we're going to talk about a few things. First, I'll introduce RegCheck, which is the name of the project that we're collaborating on. It's a machine learning based application to detect malicious registrations. Then, we'll go over how SIDN uses RegCheck in our approach, so in the pipeline from registration to whatever happens next, how do we apply RegCheck. Then, Ronald will talk about how they handle malicious registrations at .BE, and how they will use RegCheck in the future. Then, we'll discuss our collaboration in more detail, and finally, I'll wrap up with some of the lessons learned so far. And you can already go to the next one.

So, this is RegCheck. As I spoiled earlier, it's a machine learning application to detect malicious domain name registrations at the time of registration. And I presented this on Tech Day, so there is some overlap if you were there already. And if you would like more information on this, you can look back at that presentation. So, on the

schematic on screen here, you can see RegCheck consists of a couple of technical components. The first is a registration connector, which connects registrations from a registry into the RegCheck application. We have a RegCheck database that contains information on past registrations, together with labels that are used to train a machine learning model. And we have a block with RegCheck ML, which I would like to say are kind of machine learning building blocks. They're plug-and-play components, like machine learning detection methods that can be used, or features that can be taken into account, or characteristics of domain name registrations that can be looked at when trying to determine if the domain name is a malicious registration or not. RegCheck will test a domain name registration and return a risk score. So, it'll score a registration on a scale from 1 to 100, which determines or which says how likely it is to become malicious in the future. So, that's what makes up RegCheck. A registry can then use RegCheck by providing a couple of ingredients. The first is registrations, obviously. So, past registrations in combination with information from abuse feeds or other sources of labels are used by RegCheck to train the machine learning model to detect malicious registrations. Then, new registrations can be tested against RegCheck and be assigned a risk score. A registry will need abuse analysts to either interpret that risk score or to come up with a policy on how to deal with them, and some machine learning knowledge in order to fill in those machine learning building blocks that I mentioned, that last component of RegCheck. So, RegCheck is customizable to fit the registry's requirements in which features are looked at, so which characteristics are taken into account in coming to a risk score. So, a registry needs to have some idea on which characteristics to include.

Then, when RegCheck assigns a risk score to a new domain name registration and it exceeds a certain threshold, the registry will need some policy and a technical implementation on how to deal with those malicious registrations. And that policy can differ quite a lot between registries, as we'll illustrate in the following slides. Next, please.

So, the approach we take at SIDN for .nl is the following. We're running a prototype for quite a while already now. Let's start with the first thing. All new registrations enter the zone file, so we don't do deferred delegation, even for high-risk domain names. RegCheck will then retrospectively compute risk score for recently registered domain names. So, if we look at the pipeline, we start with a new registration, which is ingested into the RegCheck application. The RegCheck classifier will then look at that registration and assign a risk score to it. So, you can see this classifier as a big filter, where the input is all domain name registrations and out go only those that exceed a certain threshold for that risk score. We then publish those high-risk domain name registrations on a dashboard for our support team to have a second look at. So, they are kind of a second filter, a human in the loop that checks the results from RegCheck. And if they also think it's a malicious registration or there's something fishy going on, they will start a response, and that response is in the form of an ID verification. So, we'll send a request to the registrant to verify their identity. If they don't do that within three days, we can take the domain offline. And the benefit of this approach is that we're very accurate and diligent in our response, but the downside is that it's still relatively hands-on. We have our support colleagues looking at the results, and that the processing time is still quite long. It's three days.

This is an example of the support dashboard. So, a registration passes or fails this RegCheck test and is reported on a dashboard. And because of the simple machine learning method we use, we can interpret this risk score and assign or see which characteristics of the domain name registration contributed most to this risk score. In this case, in this toy example, it was some characters in the domain name and the address which was invalid. And in total, we see that about 44% of the ID verification requests that our support team starts, 44% of those are started at the basis of a RegCheck notification. So, we see that a lot of the work from our support team starts with RegCheck. At the same time, we see that less than 5% of registrants respond to these requests, which may indicate that we're indeed flagging those that would have become malicious in the future. I also say that quantitative evaluation here is difficult. That's because the only way to assign numbers to an evaluation here is to compare it against abuse reports, which is exactly what we're trying to prevent. So, as RegCheck becomes better, there will be fewer abuse reports for us to evaluate it against. So, in the extreme example of RegCheck working 100% perfectly, there will be no abuse reports and we would score zero on any metric. So, if anyone has any idea on how to quantitatively evaluate a system like this that we are actively meddling with the evaluation set, I'd be happy to hear it. So, that's the .nl approach. In summary, no deferred delegation. We will start with registrations afterwards and start ID verification processes for those that are high risk.

RONALD GEENS: Okay, let me take over for the next slide, please. So, in DNS Belgium, we have in fact launched a deferred delegation or delayed delegation

project already in 2020. At that point in time, we were not ready yet to start using any ML-based approach. So, we started off with a rule-based system that is in fact still currently in use, of course. The rules are evaluated continuously and adapted to whatever we see coming. So, how does that system work? A registrant registers a domain through a registrar and then the combination of contact handle and domain name is being used, is being fit into the risk-based approach, the risk-based, the risk system, in fact. And only when we assume that there is not a high risk, we will enter the domain in the zone file. If the domain is not entered in the zone file, then the registrant will get an invite to check its identity in an automated way. And only when he does that, his domain name will be entered into the zone. Sorry.

So, like I said, this is currently a rule-based system. We have been experimenting with an ML-based system on the site, and so we have been running it in parallel. But we have seen that this system has been growing and growing over quite a long period, in fact. And it actually became too complex. And then we came into contact with the work of SIDN. And now we are working together in order to, so that we can start using the reg-check mechanism to replace our rule-based system for this risk-based approach to deferring the delegation. So, a big difference already is, in fact, that we don't do delegations if we consider the risk too high. And that's already an important difference. But we also use the ML differently, in the sense that while SIDN actually optimizes to have the highest possible precision in order not to have a lot of false positives, we prefer to have a few false positives instead of being 100% accurate. This allows, the only consequence is that there is an additional ID check, but still it allows us to be more sure that we

capture more abuse, or try to capture all abuse. Another difference that we see is that SIDN uses a certain set of input data to train the system, and abuse lists typically. We have these abuse lists as well, we also use them to train our system. But on top of that, we also take into account the number of, let's say, the cases where the registrant chooses not to validate his domain name, or to verify WHOIS data. This gives us some additional data, which we hope will be giving us a more accurate view on what we really want to stop, which is mainly phishing, in our case, with the delayed delegation. Okay, that was it from me.

THIJS VAN DEN HOUT: Next slide, please. Right, so we can summarize the progress so far in three phases of collaboration. As Ron said a while ago, both of our registries started to experiment with machine learning in trying to detect malicious registrations. Then, in March of last year, we with SIDN presented at a center R&D meeting an early prototype and a feasibility study of RegCheck, and then DNS Belgium saw, hey, there's some common ground here, and we came into contact, and we started exchanging ideas and the different approaches that both of our registries take in tackling this problem. And by the end of that exploratory phase, we came to the conclusion that RegCheck is a good basis for a joint project that is the most mature and easily adaptable. So, in December of that year, we signed a collaboration agreement, and in that agreement, we established a shared ownership of the RegCheck code base and the intellectual property. And from that point on, we've been jointly developing, further developing RegCheck and merging the code base from DNS Belgium and SIDN to come to one RegCheck better version. Next slide, please.

Over the past half year or so of closely collaborating and really jointly working on this single code base, we've learned that collaborating on a project like this has quite a lot of benefits over trying to produce a system like this by yourself. And here are a few concrete examples. They're kind of technical, but I'll go through them. Most of them are things that one registry did consider before and the other didn't, and they're not necessarily implemented in the joint code base. One is improved reputation scoring. So, for some of the characteristics of a registration, we calculate the reputation over a time frame, and that resulted in some spiky results in the risk score, which we've mostly mitigated by improving the smoothing algorithm that we used for that. We've introduced the validation of geographical features like the registrant city and the time zone. We improved the way of encoding the relevance of some registrant fields using an TF-IDF-based approach, which is an information-retrieval term, like the registrant city or the registrant country. We've also changed the way that we look at the domain name as a risk factor by only taking into account those character combinations that also occur a couple of times in malicious registrations. And on a more high level, on a more strategic level, we had many discussions about the application design and the technological choices that we make in the design, which has led to a more generic and overall better code base. So, that's a big benefit. Next slide, please.

Then, to wrap up, some lessons learned over the past year. We've learned that, in the first place, sharing experiences with and different views on abuse is very valuable. So, in this exploratory phase, we were already talking a bit about what the two of our registries do to tackle

abuse and what we do to do it proactively, and some of the experiences. And that's shown to be very valuable in itself. We've learned that collaboration works even with diverging policies. As you just heard, the two approaches from .nl and .be differ quite a lot, but still we can work together on developing one single application that can be used in both of these scenarios, so you don't have to completely agree on everything to be able to collaborate.

We learned that collaboration inspires innovation, both on a technical level, so we learned new technical things from each web package we didn't use, but also on a policy level. But, of course, more people means more opinions, so the runtime is a bit longer, there are a bit more opinions on both sides about technological choices or design choices, but that's nothing we cannot overcome. And I think that's it for our slides. Thank you.

NICK WENBAN-SMITH: Brilliant. Thank you very much. So, I think that was an example of ccTLDs collaborating and working together to effectively reduce abuse, I hope. We didn't really get on to whether you think the RegCheck has reduced abuse or is it too early to tell?

THIJS VAN DEN HOUT: I think it has, because it's been increasingly difficult to find enough labels to retrain our model on purely abuse reports. So, in that case, yes, indeed.

NICK WENBAN-SMITH: Wow, okay, amazing. So, the last part of this session, and there's a good half hour allowed for it, is essentially a bit more of a dialogue on that. We'll have some Q&A time, but I think it's quite nice to have that interactive. We've also got the leadership team from, I was going to say the dark side, but actually it's our friends in the GNSO Registry Stakeholder Group. So, I've got Brian and Alan who are going to join in. So, I could be really mean because I could say you can go into the naughty step over there with all of the bad people, because it seems to me, two take-homes, firstly, that the ccTLDs, whichever way you look at it, have got lower abuse stats. And the registration process is, if you look at the friction which the Dutch and Belgians have put into this to spot high risk registrations before any harm is done from them, seems to be quite effective. I don't know if that's a correlation or a causation, and those are the sorts of things which we're quite interested to share. So, while you're thinking of your initial questions, I'll put the questions in the Q&A box. So, let's go over to Joke, who's going to read out some from Zoom.

JOKE BRAEKEN: Okay. The first question is from John McCormack. Is there a link for the paper that had parking at approximately 20% in TLDs?

SAMANEH TAJALIZADEHKHOOB: Yes. The paper is published in IEEE security and privacy European this year, and I can provide the link towards our ccNSO colleague if you like.

JOKE BRAEKEN: And then from Luc Seufer, could you explain why domain parking is listed here? Do you consider it as some sort of abuse?

SAMANEH TAJALIZADEHKHOOB: So, I talked about domain, parked domains, because in this specific session, because that's -- when we talk about the domain space that is included in the zone files, we consider that all of the domains are active domains, as well as domains that are listed on the reputation block lists. Now, if you take the domains of those lists of every day and check whether they are active or not, you come to a different conclusion, and one of the categories in which domains can be inactive is parked. There are many more, at least in the way that we categorize in our research, and hence the research that was published this year. Thanks.

NICK WENBAN-SMITH: Are there any hands or questions in the room? Mr. Disspain.

CHRIS DISSPAIN: I've just got a couple of questions, really mainly for clarification, I think. I think Nathan, as a start, Nathan, I may have misheard you, but I think on your slide you talked about DNS abuse, but then you said, and abuse on the Internet. And I just wanted to figure out whether that was actually a wider thing or do you just mean DNS abuse. And secondly, I noticed that you -- where you're pulling some of your data from, and I wondered if you have a stated definition of what you think DNS abuse is. Do you subscribe to the same definition that COMPASS uses, et cetera, etc.? Thanks.

NATHAN ALAN: Yep. So the feeds that we used, the sources of the abuse that we measured were from the likes of URL house, OpenPhish, APWG. We did include SpamHaus spam data in this study. I will say that the data that we have available is on the website, and we are looking to improve that so that you can filter out certain types of abuse if that's what you would like to do. But, yeah, I think the way that we define abuse is very similar and probably identical in many ways. And, yeah, it was just the abuse from those feeds, the DNS abuse, yeah.

NICK WENBAN-SMITH: One more remote question, then we'll go to the hand over there. One second.

BART BOSWINKEL: Hi, Nick. There was a question from Luc Seufer ... around RegCheck. He was wondering whether he missed the false positive percentage, or it wasn't mentioned. So this is more for the SIDN and DNS BE Corporation. Thanks.

THIJS VAN DEN HOUT: The false positive percentage, it's been fluctuating a little bit. It depends on what you call a false positive. If we compare our results to Netcraft results, historical Netcraft results, so abuse reports, we achieve about 20% precision, I believe. But it's difficult to evaluate because RegCheck finds things that Netcraft finds and doesn't find, and Netcraft finds things that RegCheck doesn't find. So they're kind of complementary.

It's very difficult to see how many false positives are truly false positives or if they are just finding things that Netcraft didn't find.

Overall, on the abuse dashboard, so the things that Netcraft posts on the dashboard for our support colleagues, I think about 25 to 50% are followed up on. I'm not exactly sure about the numbers currently, so I'll have to leave it at that.

NICK WENBAN-SMITH: Great, thanks. May I have a hand over there, please?

JUTTA CROLL: Yes, thank you for giving me the floor. From these presentations and also from the previous session at the GAC, I understand the definition of domain name abuse that is a correct at ICANN, but still I'm wondering when per definition spreading malware is considered domain name abuse, but when certain domains are used to spread child sexual abuse material, that is not considered to be domain name abuse, and I'm wondering whether this could be debatable or whether it's not a topic for an ICANN debate. Thank you.

BRIAN CIMBOLIC: I just also want to correct that Alan and I are from the gTLD Registry Stakeholder Group, not the GNSO in particular, but it's a great question, and I think that when we're talking about DNS abuse, people often conflate just because something is not DNS abuse doesn't mean that it's not terrible, and it doesn't mean that it shouldn't be addressed online, but when you hear registries in particular talk about DNS abuse,

what we're really talking about is abuses that are appropriately handled at the DNS level, and so child sexual abuse material is obviously terrible, and I know I and Identity Digital and Nominet, we each have robust programs to deal with child sexual abuse material, so we do take action on it, but just because we take action on it does not mean that it's DNS abuse, and I know that sounds confusing, but if you think about it, 99% of the child sexual abuse materials that we come across are on file sharing sites, and so what that means is that you might have a percentage of a percentage of a percentage, so .00 something of a website's content is child sexual abuse material, but for us at the domain name level, the only thing we can do is suspend the entire domain name, so you're knocking down access to legitimate content, it may have nothing to do with child sexual abuse material, and so because we have such a limited tool to deal with this, we have to be very careful when we're talking about what is appropriate to use the DNS to remediate, and so when we say DNS abuse, it doesn't mean that everything else isn't terrible, it just means that generally speaking, we're not the right actors to deal with it because of the collateral damage that can occur when we take that action.

JUTTA CROLL: Thank you so much for your answer. We have different data from the [hotlines] and the InHope network that it's moving from the file sharing service to sharing on certain domains, so probably it's good to look into that further on. Thank you.

ALAN WOODS: This is actually just going back to maybe Chris Disspain because in fairness, Brian just answered that I think exceptionally well, so I have nothing more to add there, but going back to Chris Disspain's question, specifically as well when we're talking about -- actually it might not have been Chris's, but it was a question about sources, and I think this is one of the perspectives that I can bring from our experience within the registry stakeholder group and the ICANN context. For many years, we have discussed the concept of DNS abuse and the concept of what sources are those sources that should be used, and we have really crystallized around a concept that any source is a good source as long as that source is capable of being evidenced. So some of the sources that are mentioned up there have traditionally not necessarily been evidenceable to the extent that we as registry operators can rely on them in order to take the actions that some expect us to do. So a lot of the work over the last eight years has literally been in bringing the conversation towards that. As long as there is evidence and it is appropriate for us to take action based on that evidence, then that is something which should be something that the registries and registrars can look into, and perhaps where appropriate, take direct action or take other actions. And I think it's very important as we're looking at the tools for measuring DNS abuse or abuse online, that we take into account that core principle of transparency and non-arbitrariness and making sure that we give credit to those operators who are looking for those elements of evidence, expectations of transparency, and indeed when actions are taken that it is auditable, that you're able to trace that action, why you are taking that action, based on what evidence, and indeed there can be inferences coming from that evidence, but we need to be sure that we are not ignoring due process. A report of abuse is not

t he same as an evidence of abuse and we need to be very, very, very clear on that.

NICK WENBAN-SMITH: Thanks, Alan. If there's no more questions in the chat, I've got a few questions of my own and I'd like to ask all the panelists really. We seem to be getting a handle on measuring, perhaps early stages. I see that whether it's ICANN, DAAR, whether it's the COMPASS, whether it's the Research Federation, they all have quite a distinct difference between gTLDs and ccTLDs. I'm interested to understand that a bit more because it doesn't seem to be exclusively on size or registration policy. I have to do this because I can see Andreas and we're in beautiful Hamburg. The German ccTLD, the .DE, is very open, it's a very large zone, and it is amazingly well run and safe. In some ways, I'd say .UK is similar. It's more like, in terms of open registration policy, low price, it's more like many of the gTLDs. But the observed abuse, it seems, on three different metrics, so they may not all agree, but the trend seems very clear and it's not just by 5 or 10%, this is by orders of magnitude. Why is this? Can anyone explain it to me? Because I've been thinking about this quite a lot and I'd like to know the answer, please.

ROWENA SCHOO: When we were putting together our public reporting, we were thinking about how do we construct this and what groups do we put TLDs into, and that was a really challenging endeavor. We eventually just split them into gTLDs and ccTLDs. But when I was writing the explanatory text for each group, it occurred to me how much variation there was within both groups. We often talk about variation in ccTLDs, but there's

also a variation within gTLDs. Some of them operate in a very closed model. For example, things like .bank, .pharmacy, and that naturally means they have a restriction on what comes in, which tends to mean they have a restriction on abuse as well. In some of the ccTLDs, we see very restrictive models that probably are not generalizable to the wider open ccTLDs. For example, we know our Norwegian friends have a very restrictive model that reduces what comes in. In terms of pinpointing exact reasons why there's more or less abuse or the kind of drivers of that, I think I would actually point to an interesting study that is taking place at the moment that ICANN is working on with Maciej that is going to look at these exact factors. How do different price points and different issues influence how abuse is driven? I feel like you had more questions that I haven't answered, but maybe that's enough to start people thinking.

NATHAN ALAN: In the study that we found, it just seemed like there were more checks in place. Whether that was pre-registration, after registration, understanding their customers, it was a combination of different things that seemed to mount in the lower levels of abuse. That was just something that we had observed from trying to figure out, or having the registries inform us what they had been doing in terms of registration points.

NICK WENBAN-SMITH: Thanks. Anyone? Samaneh?

SAMANEH TAJALIZADEHKHOOB: To just complement what the others already mentioned, I think, like Rowena said, the main thing is that you need to find a factor to attribute, to categorize, for the sake of presentation. But also, the general rule is that the attack space, which is here the number of domains, is the most critical determining factor in the number of reports and eventually incidents. And that is just the size of TLD, nothing else.

NICK WENBAN-SMITH: Thanks. Are there any other thoughts? Alan?

ALAN WOODS: So, two anecdotes. The first one was, and this is not going to help, but it will put thoughts into people's brains, that recently I had experience of a specific attack, a DNS abuse attack, whatever you want to call it, was a campaign that focused on a relatively quite expensive domain. So, most people in these things, they go to, oh, the cheaper the domain, the more likely it is to be bought. But in this instance, we saw thousands of domains being registered at a price point that was anywhere between 30 to 60, because the benefit that was being derived from the domains could potentially have that have been a drop in the ocean. And that's the other issue, is that methods are much more sophisticated, so simple metrics that we all blame might not necessarily be the ones that we expect anymore. Now, that being said, I'm sure that there's still a factor very much so. However, as we're having much more mature conversations in this space, we're finding that not always does the rule follow, and that I did have a second one, but of course, in the grand scheme of ICANN, cannot remember what it was, so I'll stop there.

NICK WENBAN-SMITH: Yeah, no problem. The price is interesting, and it did come up this morning in the GAC presentation. There's been a long-held anecdotal assumption that cheaper domains equals more abuse, but there's plenty of good examples where that's simply not true. So you need to be quite careful, and it's a multi-factor thing which leads into this. So, Rowena.

ROWENA SCHOO: I remembered a few more things. I think I mentioned in my presentation as well a point about geographical bias in the list as a potential. Most of these sources that are used target global brands, and that might not be representative of all countries and all local targeted entities. I know from some of our conversations with other countries, they find phishing attacks that aren't present on these lists, but a lot of them operate a model where they report that to an ISP blocking scheme. They feel they've protected their citizens. Those types of attacks might not float to the top. There is also potentially a scaling of cybercrime, too. Do you try to target as many people as possible? Is it more helpful to do that if you're using a generic top-level domain that might be recognizable to people all around the world? I should say, these are thoughts, things that we're discussing. I don't know if there's research to specifically pull that out as conclusive yet.

NICK WENBAN-SMITH: As I heard that, I think what you're trying to say is not so much that the people in Germany are more honest, it's just that sometimes people use

generic top-level domains because the addressable market for targeting abuse is bigger. I understand that. Bart, I see you have your hand up.

BART BOSWINKEL: This is again a comment from John McCormack from Hosterstats. This is a comment on one of your questions. One potential answer has to do with trust. The ccTLDs seem to be higher trust TLDs than gTLDs, which have global markets. People identify with their local TLD as being their TLD, but the effect is not as common with gTLDs.

NICK WENBAN-SMITH: Thanks, John, for that comment. I just want to challenge a couple of those points. It is true that some gTLDs are genuinely gTLDs and generic in that sense, but if it was the case that a generic type of name-- because we all know that there are some ccTLDs which are also generic or quasi-generic, whether it's-- I'm just picking examples, but .co, .me, .tv, there's lots of-- even .it for the Italians has some sort of generic thoughts. You would have thought that those names would perhaps follow a gTLD pattern or ccTLD pattern, but I don't believe that there is any evidence for that.

ROWENA SCHOO: Thanks, Nick. I think that's a really interesting point, and one of the things we are starting to observe through COMPASS when we look at malicious registrations and compromise is that there is a little bit of a pattern appearing around the ccTLDs and the older generic TLDs having more compromise. So more domains that are registered for otherwise

legitimate benign use that have usually become-- they've had a content management system, it hasn't been updated, there's a vulnerability, it's become compromised. We tend to see more of that in ccTLDs and also the older, larger generic TLDs, the original ones, whereas when we see newer gTLDs that are growing at quite a rate, they tend to have a different split in the type of abuse where they have more malicious rather than compromise for having less domains for a less period of time.

NICK WENBAN-SMITH: Thanks. Actually, I think I heard that right. It's not that there's less abuse necessarily in ccTLDs, it's just that it's a different type of abuse. It's not necessarily malicious registration, it's just that you've got a lot of existing registrations and lots of old websites to which these domains direct, and so the abuse is not-- or perhaps criminals target websites with vulnerabilities rather than exploiting weaknesses in the registration system. And perhaps-- I suppose this is what has come out of a lot of the tools and management discussions, which is you think you're comparing gTLDs and ccTLDs, but actually the numbers are quite skewed because you've got a very large number of new gTLDs, so that changes it. So you're not necessarily comparing the same types of attributes and cohorts when you might have thought you were. So I think that is an interesting point, and I think for us in the established, mature ccTLDs, we would be doing better if we could find ways to contact the registrants whose websites have been compromised and get them to fix their vulnerabilities as opposed to making it harder to register new domain names. Anyway, so that's just my reflection there.

We talked a lot about reputation block lists earlier, and I wanted to know-- because there's so much choice, and somebody-- so because I'm stupid and lazy, I want to copy somebody else's homework. So you've all looked into these reputation block lists. Which one-- if I've got one to buy, which one should it be?

ALAN WOODS: I'm not going to give you a straight answer on that because, you know, at the end of the day, I am a lawyer. But to be perfectly honest, this is one of the things that I don't think we engage enough in in our industry, and the answer to that is we also need to be very mindful that the quality in different block lists can also, over time, vary. So we should be not only measuring the abuse but also how good or how consistent or how transparent or how beneficial the block lists that we use are. I'm just going to go back to what I said earlier and say I would use any list as long as that list is capable of grounding an evidence-based escalation. It's as simple as that. I don't care if it's a single report or a thousand reports that come in, as long as the quality and the elements in that report are showing that there is an abuse occurring and that we are the appropriate party to take action.

NICK WENBAN-SMITH: I think you've just described a unicorn, by the way. I've not yet found one which does that. Rowena, I think you have your hand up.

ROWENA SCHOO: Thanks, Nick. I was more going to comment on what you were saying before, which is just a caveat that there's a lot of variation in abuse rates

in ccTLDs in the same way that there's a lot of variation in gTLDs. We publish reports on this. We have tables on this. Those figures move around, and sometimes there are ccTLDs that do have high rates of abuse, whether that's malicious or compromised. Come and see us and see your dashboards to find out if you're one of them.

NICK WENBAN-SMITH: Very clever to get a little bit of self-promotion in there. I have Brian and then Bart, but if there's anyone else in the room who wants to ask a question, if they could catch my eye. Thanks.

BRIAN CIMBOLIC: Thanks, Nick. I want to pick up on something that Alan noted, and that's reputation block lists are important and are an important indicator for further investigation. It would be a mistake to rely, you know, sort of sight unseen on reputation block lists alone. I think Alan's point on evidence is really important. I'll give you a couple of concrete examples. I've had reputation block list providers give me lists of domains for us to take action on to include not only domains that weren't currently registered but had never been delegated before. I've had reputation block lists -- this one's not a joke -- someone gave me a list of domains and included on them was the RBL provider's own domain name. And so there's varying degrees of quality that go into this, and that's not even for me to necessarily say that these things are bad, but they're oftentimes used for different purposes. They're used for network blocking. They're used for other sorts of things that it's just usually the tool isn't fit for purpose when just applied one-to-one on the domain

name level. So, again, that's not to say reputation block lists are bad, but very few of them are really tailored for use at the DNS level.

NICK WENBAN-SMITH: Thanks. Samaneh.

SAMANEH TAJALIZADEHKHOOB: To continue what Brian mentioned, we at OCTO have done a study on this, and we published, if I'm not mistaken, five to ten metrics, things that are measurable for block lists to evaluate how good they are depending on the purpose of use and things that are harder to measure. This paper is published in this year's e-crime conference, and we can provide the link to that.

NICK WENBAN-SMITH: Thank you. I mean, there's a few more minutes, and I've got a few more questions, but I'm interested in future work of the DNS Abuse Standing Committee for the ccNSO. On Saturday, I did set out a future work plan, and one of the things that we were considering for a future meeting would be a bit more of a practical workshop around your own ccTLDs, tools and measurements, and a bit more hands-on how to. I'd be interested to know whether the people here at least indicate that that's something useful, that there's more interest in following up on this session of tools and measurements in terms of that sort of thing. I'd be interested if people either message me privately or wave at me or something, if people are interested to do that, then it's something that the leadership team of this group would look at going forwards. So any thoughts on that sort of area, or do you think we've solved all of the

metrics and measurements problems? Or none of the above, potentially. So we just have got another five minutes. Is there any more questions in the room? Andrew. I'm

ANDREW BENNETT: I noticed earlier that you said the DAR was opt-in and accounted 22 ccTLDs on there. I think there's over 300 ccTLDs. I just wondered what was the reasons that some of the larger ones don't opt into DAAR? Is ICANN able to persuade more ccTLDs to opt into that system?

NICK WENBAN-SMITH: Samaneh, you can answer that. I also have some thoughts on that.

SAMANEH TAJALIZADEHKHOOB: Actually, the short answer is that we don't know. We are putting effort in bringing on forward how the project is and what are the benefits. We, over time, learned that maybe some ccTLDs have concerns about their zone files being shared or used, which we repeated multiple times that this is not going to be the case. The zone file is only used for this project and the reports are sent to the admin contacts that we have from the ccTLD. And the trend over time as the project leads, we observed, is that it's gradual, but it's growing. And we are really engaging with the participants of the project to receive their feedback, make the reports better and improve over time. And we hope that it helps. But again, this is just where it can go so far.

NICK WENBAN-SMITH: I think that's exactly right. There's a sensitivity with some ccTLDs over handing over the zone file and for them it's a hard no. And I think from others I've spoken to, you get your dashboard but you don't get any actionable data from it. So there's a sort of a, well, there are some technical requirements and things to go into to set it up and do you get back enough of a return, I think, has been the question. But I see there are some comments in the chat. Is that what you wanted to read out?

BART BOWSINKEL: Yes. Two comments. One again from John McCormack from HosterStats. This goes back to your earlier question to the panel. Panelists, is the registration lifetime of a lot of new gTLD domains is lower than those of legacy and ccTLD domain names? This high non-renewal rate limits some of the possible compromises to software like WordPress plugins. This can skew the figures. So that was the comment from John McCormack. And I have a comment from Alexander Hugla from Gandhi. Private reputation lists funded by private companies are unreliable because they have a bias towards their clients and their client advocacy. These were the final ones.

NICK WENBAN-SMITH: Excellent. Thanks, Bart. And it's Nick again. Yeah, I must admit I have seen some reputation block lists which are marketed on the volume of their names blocked rather than the quality of the diligence and evidence which goes into the blocking decisions. So I can understand that comment. So thank you. I just wanted to show just in the last minute that the work of the DNS standing group does not stand still. We are thinking forward to the next couple of ICANN meetings in terms of -

- so in the Cancun meeting when we first presented the survey data, I somewhat unwisely asked for suggestions as to what was worthy of further research and discussion. And these are the four suggestions from ICANN 77. You can see the third bullet point largely matches the main topic of today's thing. So we've tried to reflect back to the community what the community asked us to do. So it's very nice to see such a well-attended session and to have such good engagement from everybody. I'd like to thank everybody for that. You'll see there are three other topics, and we are in the leadership team looking at scheduling that over the next six to 12 months so there's a forward-looking workflow of other things which people have asked for. I would say that this is not set in stone. We're very responsive to requests if there's anything that people want to add to the list or if people don't think that they are so interested in any of these topics anymore, then we can definitely modify that and you can contact any one of us. So that's just to give people an idea as to the forward workflow on this. I'm particularly interested in examples of good collaborations as between different registries so we don't shift abuse from one registry to another, whether that's a gTLD or a ccTLD. We should try to do this in a coordinated way just to shift it away from TLDs. I think it would be quite nice also to see some good examples of how we can more effectively work with registrars. There are some registrars in the room which is great to see, but how can we do this better because collectively I think we have a lot more information and resources to act on than we each do individually. So that's probably going to be one of the main topics in the next ICANN meeting in Puerto Rico.

ROWENA SCHOO: I just had one final plea for information because talking about lessons that can be understood and learned about different registry operations, I'm personally really interested in whether incentive schemes have an impact on registries being able to manage their registrar channels to have low abuse rates. I'd like to look at that, but I'm not sure if I know all of the TLDs that run an incentive scheme. So if you run one or know of them, I'd love to hear about it. Thank you.

NICK WENBAN-SMITH: Okay. With that, I'm going to wrap up this session. Thank you very much for everybody contributing, and see you all again soon sometime. Thank you very much for all my panelists and contributors.

[END OF TRANSCRIPTION]