YEŞIM SAĞLAM: Good morning, good afternoon, and good evening to everyone. Welcome to At-Large Consolidated Policy Working Group call taking place on Wednesday, 17th of May 2023 at 13:00 UTC.

We will not be doing the roll call due to the increased number of attendees as well as for the sake of time. However, all attendees both on the Zoom Room and on the phone bridge will be recorded after the call. To cover our apologies, we have received apologies from Maureen Hilyard, Steinar Grøtterød, Mouloud Khelif, Judith Hellerstein, Cheryl Langdon-Orr, Claire Craig, Greg Shatan, and from Chantelle Doerksen. From staff side, we have Heidi Ullrich and myself present on today's call, and I will also be doing call management. As usual, we have Spanish and French interpretation provided. We have Claudia and Marina on Spanish channel, and we have Dominique and Isabelle on the French channel.

Two more reminders before we start. The first one is for the real-time transcription service. I'm sharing the link here with you. Please do check the service. The final reminder is to please state your name before speaking, not only for the transcription but also for the interpretation purposes, please. With this, I would like to leave the floor back over to Olivier. Thank you very much.

OLIVIER CRÉPIN-LEBLOND: Thank you very much, Yeşim, for this introduction. Welcome to this week's Consolidated Policy Working Group call which is going to consist mostly of a visit and a presentation from Graeme Bunton, who is the director of the DNS Abuse Institute. He's going to provide us with a full

demo of the NetBeacon. So we look forward to seeing this. That will take about 40 minutes of our time.

After that, we will be focusing on the workgroup and small team updates with just an update from the Expedited Policy Development Process on the Internationalized Domain Names, the EPDP on IDNs. And one very short update on the Applicant Support GNSO Guidance Process. After this, we'll have our usual policy comment updates with an update on the policy pipeline that is taking place. And you'll notice that the main one is the proposed renewal of the Registry Agreement for .NET, which will be dealt with in agenda item number six by Michael Palage and Bill Jouris. After that, we'll have Hadia Elminiawi and Jonathan Zuck taking us to ICANN77.

Is there any other business? Are there any other items to add to the agenda? Is there anything to be shuffled around in case you need to leave very quickly and very early, or shuffled later in case you're not here yet? In which case, you probably are not hearing what I'm telling you. I am not seeing any hands up. So it looks like there's a question. "Is there a link to that draft African DNS report?" I do not know. It will have to be checked. Anyone can please respond in the chat? I'm not seeing any hands up. So the agenda is adopted as it currently is listed on your screen.

We can look at last week's action items, of which there's one remaining which is for Michael Palage and Bill Jouris to finalize a draft statement on the Public Comment proceeding on the proposed renewal of the Registry Agreement for .NET. That's going to take place during this call.

Everything else is done. Are there any comments and questions on any of these action items? I'm not seeing any hands up. Okay. Thank you.

I note in the chat that Heidi is saying that staff will note the request as an action item, the request being finding the link to the draft African DNS report. Thank you, Heidi, for this. The action items are completed and checked.

So now we come to our star of the day, Graeme Bunton, welcome, director of the DNS Abuse Institute. It's not the first time that Graeme comes on one of our calls and visits us and explains what he's up to and what the DNS Abuse Institute is up to, but it is the first time we're going to get a NetBeacon demo. So without further ado, over to you, Graeme.

GRAEME BUNTON: Thank you. Now, we've got a little bit of high jinks in place today because I'm going to push my video through Jonathan so that we can record it. But we'll start that in a minute or two. And so maybe I'll do a brief intro. Is that okay, Jonathan?

JONATHAN ZUCK: Yeah, that's fine.

GRAEME BUNTON: Okay. So the way this is going to work is I'm going to switch screens and share my screen, and we'll push that video through Jonathan so that he can record it, we can make a video and distribute that video so everybody can benefit from this demo. I'd like to reserve a bit of time at

the end to do some Q&A. We can talk through some pieces in more detail, and probably also a little bit of time to talk about other activities in the DNS Abuse Institute. But briefly, as a sort of way of introduction before we get into that bit, my name is Graeme. I'm the director of this project. The Institute is a project of Public Interest Registry, who operates the .ORG TLD, and is a not-for-profit. So as part of their not-for-profit mission, they were like, "We can do something about malicious domain names on the Internet. How do we do that? Let's start this thing and see if we can come up with some new ways to solve abusive problem." So we're going to talk today primarily about one of those. We have some other projects I'll talk about shortly after.

I'm coming to you today live from sunny Barcelona, although normally home is Toronto, Canada. I'm here trying to do a whole bunch of work across the European continent and it just made more sense than flying back and forth across the Atlantic a bunch of times. It's very nice place to be.

So the demo, the actual work itself, I think will be reasonably short. It's pretty straightforward very intentionally. We want this tool to be easy to use for everyone. So I think where we'll get into that full 40 minutes, it is probably where we have some back and forth some questions. How do we think about expanding this tool? Where are we going from here? So I think with that, I'll get into it. I'm going to jump over to this other thing and push my video to Jonathan. Are you ready to go, JZ?

JONATHAN ZUCK: I'm ready. Yeşim, showcase my video if you would. There we go. And then you need to share your screen, Graeme.

GRAEME BUNTON: Great. So this is going to be a demo of a service that we call NetBeacon. NetBeacon is a free for anyone to use, centralized abuse reporting system. We created this service primarily because reporting abuse online is complicated. It has a sort of technical barrier to entry that is beyond many end users. So we've tried to remove that technical barrier. On the other side of abuse reporting is the reports that registries and registrars get, and they have for a very long time been problematic. They're duplicative, they're unstructured, they're unevidenced. And so what we've tried to do by building this essentially abuse reporting intermediary is to simplify the process for people submitting abuse, as well as to improve the results of those submissions to registries and registrars so that they can do something about that abuse.

NetBeacon is available at netbeacon.org. Anyone can go there. The actual website of netbeacon.org is mostly explanatory. The pool itself is at app.netbeacon.org, but you can just click this big Report Abuse button and it will take you there. And so when you get to the actual Abuse Reporting page, it looks like this. You'll note the first thing is that you can use single sign-on with Google. We'll look at integrating other single sign-on partners, ultimately. Google was sort of the easiest, biggest one to integrate to start, but we would love to do Apple. We were considering Twitter but that got weird, potentially Facebook. Or you can just create an account using your own e-mail address and

password. I've already done that for the sake of simplicity of this demo, but you can do it.

It is worth noting here, though, that you need to have a verified e-mail to use the service. Verified, meaning you've signed up and clicked the link that's landed in your inbox, which means that there's no anonymous abuse reporting going into NetBeacon. And that's important for a couple of reasons. One is that if you submit an abuse report and the provider, the registrar primarily, requires more information. They need to be able to contact you to ask for that. The other is that abuse reports are frequently weaponized. So sometimes people use abuse reports to silent speech to try and remove competitors from the Internet. So by asking people to have an e-mail address, it's a relatively simple and light way of having some accountability. So that we, as operators of NetBeacon, get notified that somebody is abusing the system, we can limit access or constrain their abuse reports that are going through.

So I'm going to log in. I've got a password—and you can tell that I've got a boatload of passwords probably for this service that I test it all the time. I'll admit I'm zoomed in a little bit here. So the UI is normally a little bit wider. I wanted to make sure everybody could read. Then also the video, I think, might be a little bit fuzzy but that's because we're recording this for posterity.

So it starts with a relatively simple form, which is we want to put in the URL of the domain name that we're reporting. So I am using an example today of a phishing attempt that I personally received and submitted to NetBeacon some time ago, originally. So we're not actually going to

submit this. But here's what I got. I got an SMS to my phone, it says, "RBC important message from Royal Bank. Please visit royalbkalerts.com." Now, Royal Bank is a large Canadian bank. I know the brand, but I'm not actually a customer. So that, to me, was a very good sign that it was a phish, that boy, an awful lot of Canadians are Royal Bank customers. So I took a screenshot of this SMS that's got the domain name in it. I also took a screenshot of the page that that pointed to. So I could say, "Hey, look, this has got a Royal Bank login, and it says my account has been suspended. I don't have an account with Royal Bank. So this is, to me, very clearly and obviously a phish." And now I've got nice good screenshots of what was happening.

A note here, you need to be extremely cautious about resolving those URLs that people are asking you to go to. And in general, you should not do it. You should leave that for professionals. So usually, sufficiently, the screenshot with the domain name in it will be sufficient.

So I now have the domain name, royalbkalerts.com. I'm going to put that in the address here. I'm going to type it really badly because everybody's watching. So the form recognized it's a domain name, I can hit Continue. And I can see that this was registered with Tucows. A note that Tucows reacted to this abuse report when I initially submitted it very quickly. Kudos to them.

So now I get to choose the type of harm that I think I'm reporting. I will say, ultimately, we might work on this UI a little bit, maybe do a more of a question-based approach, make it even easier for end users, because not everybody is capable of identifying what type of harm it is. But in this case, as it stands right now, we're asking you to identify what type

of harm it is. NetBeacon accepts abuse reports for these four types: malware, phishing, botnets, and spam. Most of what we see going through these days is phishing. A little bit of malware, a little bit of spam.

So if you click Report Phishing, it's going to tell you what information is required for that abuse report and what information is helpful. It's important to note that without the required information, a registrar is going to be unable to act on that abuse report. So submitting an abuse report without that information is actually counterproductive because abuse report management is a zero-sum game. If a registrar is dealing with one that they can't actually action, they're probably not getting to another that they could. So we want to make sure that every abuse report that we submit to this has as much information as possible. Let's hit Continue.

This brings us to a new form. It's a phishing abuse report for royalbkalerts.com. Let's step through these pieces. So when did this harm happen? We can say today, even though I'm using a historical example. We want to say which institution is being targeted. In this case, it was the Royal Bank of Canada. I don't actually know offhand what the Royal Bank domain name is. But let's guess that it's that. Here we can provide—actually, I should back up. That was maybe a little bit quick.

What I'm trying to put in here is give the registrar the information on who is being impersonated in this phishing attempt. So not me but the Royal Bank of Canada. If you see it's a fake Apple login, you would want people to know that it's an Apple. Not every registrar knows every

brand, and so sometimes these things aren't obvious. This domain was registered at a Canadian registrar. So probably the Abuse team at Tucows is going to know who Royal Bank is. But it could totally have been a registrar in France, and they may not know who that bank is. So helping them understand who is being targeted by that phish is helpful.

Here we would provide a brief description of the issue, which, in this case, would be something like "Received an SMS phish for the Royal…" That's probably sufficient in this case. Getting this right is a little bit interesting. Registrars don't want to have to sift through an essay. They're looking for simple facts so that they can get to as many abuse cases as they can.

Here we can provide additional evidence. I have those two screenshots saved. So it's pretty easy for me to go attach them. Then I can hit Continue. It's worth noting that I can save this at any time so that I can come back to this abuse report, so that if you don't have everything ready, you can save it and come back at any time.

We also ask for location. We do that specifically because some online harms are what's called geofenced, and that they're only available to people within a certain geography. Typically, that's done by IP address. Sometimes it's done by what's called the browser user agent as well. So that unless you're on a mobile phone, using a particular carrier, an IP address from a particular geography, you can't see that harm. So as a registrar is trying to verify this, it's important that they have some understanding where the geography is.

In theory, I can click this button and it might do a smart thing where it says, “Hey, that’s where I am.” That’s great. If this phish came from an e-mail, we could provide the sender’s address. In this case, we’re not reporting an abuse that came via e-mail. We also have room to capture the e-mail headers and body if this again was an e-mail-related phish. A note about this, which is that reporting abuse that’s only in an e-mail and not at a resolvable website, it’s extremely important to be able to provide e-mail headers. And not everybody knows what e-mail headers are. Typically, in your e-mail client, you need to go to something like view e-mail source, something like that, view original, and it’s going to provide you with a bunch of gobbledygook-y looking code. That code is really important for registrars to be able to figure out exactly who sent this e-mail. Then by looking at the body, they can figure out exactly what domain names and URLs are involved in the phish. So really important that you’re able to get that if the abuse is e-mail related. In this case, that’s not.

Then I would click Submit. And that’s pretty much it. It would send this abuse report to the registrar. It would not only include the screenshots that I’ve attached but NetBeacon is going to try and get screenshots itself. It’s also going to take that domain name and it’s going to check it against the number of online sources of abuse information. So that’s going to be things like Google Safe Browsing. It’s also going to check against the large reputation block list providers, Spamhaus and SURBL, as they’ve donated their access to their lists. We check a service called CrowdStrike. We check another one called URLhaus and a couple more. Essentially, what we’re trying to do is provide a registrar with the most information that we possibly can.

Speaking of donations, I would be remiss if I didn't mention that NetBeacon was built by CleanDNS. They donated both some of their underlying technology they use in their service, as well as a bunch of dev time to customize it for these purposes. So a big thank you goes out to them for that.

So that's how you would report phishing. We can go back to report abuse. If we went back, we can report other types of harms. They're all very similar in that you're putting in the domain name, you're selecting the type of abuse, and you're stepping through a similar form. For malware, we're looking for the malicious file name. We need to know where it's located, screenshots. Location again for malware, if you have it, again helpful are MD5 checksums. These are signatures of malware files. Not necessarily most people don't have access to that sort of thing. But I think you kind of get the premise from that: simple, easy-to-use forms, relatively straightforward evidence requirements, and without having to have the burden of understanding or being able to understand how to locate the registry or registrar to send that abuse.

A couple other caveats that are worth noting. Right now we're working with only generic TLD registries and registrars. We can report abuse to all gTLD registrars, that's because they're obligated to have an abuse contact under their contracts with ICANN. So they don't need to sign up or create an account, we can report to everybody. So far, that's going pretty well. Registrars have more features. They can create special accounts with NetBeacon so they can customize how they get reports. But that's not necessary for this audience. We are working on integrating ccTLDs. It's complicated as ccTLDs are all different in some fashion. But I will hope to see those begin to come online probably

within the next few months. We'll be starting with the sort of larger, easier to work with ccTLDs, and then working our way through that list.

Is there anything else that I think I need to show you? It's probably worth showing you that there are settings. I can change my identity in here. I have been verified. You can request verification. So requesting verification is for abuse reporters who do this—professionally is maybe a strong word—but with some regularity and with some competence. By requesting verification, that gives you access to settings for outbound reports, which allow you to include your organization name in the title of the abuse report that a registrar would get so that they would know who it's coming from more obviously. And then it also allows you to include some boilerplate text at the top of each abuse report that comes from NetBeacon. So maybe you're a cybersecurity organization and you would put your organization name in the top. And in the report header, you could say, "Here's our organization, here's our website, here's our contact information, if you need to know something more about us." So that feature is available to regular reporters. You can request verification in your settings.

The last thing I'll point out is that there is an API. Again, if you're doing this at some sort of scale and you have competency to do this, you can submit abuse reports directly into the NetBeacon API.

I think that's all that I was really interested in demoing here today. I'm always happy for feedback. I can see that there's a bunch of questions I haven't been able to follow in the chat. But maybe what I'll do is stop sharing my screen from here and we can go and address those questions.

OLIVIER CRÉPIN-LEBLOND: Thank you very much for this, Graeme. We see Jonathan Zuck now, star of the day. Jonathan, did you want to run this part of the call and moderate this thing? He has a problem in his system that is not allowing—

JONATHAN ZUCK: There we go.

OLIVIER CRÉPIN-LEBLOND: Okay. Here you go. Okay.

JONATHAN ZUCK: Can you hear me now?

OLIVIER CRÉPIN-LEBLOND: Over to you, Jonathan. Yeah, you can moderate this.

JONATHAN ZUCK: Okay, happy to moderate it. I've got just a quick question for you, Graeme. In order to demo this, you were using an old report or something like that. If one of us was trying to demo it to somebody, is there like a suggested way of doing that so that we're not sending a duplicate report? I'm trying to think of what's the best way. Is there a fake report that you use a lot? What's the best way to give a demonstration if one of us is trying to do so?

GRAEME BUNTON: Oh, that's great. Boy, we're also happy to do demo. So you don't have to do that. I'd appreciate it if you do. The easiest way I think is to just pick a piece of spam out of your spam folder, that almost I'm sure everybody here has some spam. So it's a little bit less interesting than phishing, which is sort of one of the reasons I didn't use it here. But pick a piece of spam. My caveat to that is that a lot of spam comes from Google, Gmail addresses, and reporting spam to the Gmail registrar, which is I'm pretty sure Mark Monitor probably isn't helpful. Mark Monitor is not going to take gmail.com offline because of spam. So probably finding a piece of spam that is not from Gmail or one of the super large e-mail service providers will be more effective.

JONATHAN ZUCK: All right. Thanks. I'll follow up with you. It's just we have the plans to maybe try to put together presentations down at the lower levels of our infrastructure. The ALSes may be giving a demo or something like that, and we'd want to write that and construct it and figure out the best way to do that. But I'll get back in touch with you on that. So I will call on Sébastien Bachollet.

SÉBASTIEN BACHOLLET: Thank you very much, Jonathan, and thank you for the presentation. I have one question and one comment. I'll start with the comment. I tried to explain to people when they get SMS, don't touch the link and don't try anything. That means that you are saying the reverse. You're saying go and put the information. Can't you find a way to allow us end user to

send directly the SMS to somewhere to take care of that? Because in general, I feel that even if you try to do it, not so complicated, it's still very complicated. For normal end user, it will be too complicated. Therefore, that means that there are only some specialists who will be able to take care of that and not doing and the other will not do anything on that on the question with DNS abuse. Thank you.

GRAEME BUNTON: Thank you, Sébastien. So I was not trying to say the opposite. Generally, your advice is very correct, do not click those links, do not go to those URLs. It should be perfectly safe to take a screenshot of an SMS that includes a link. So that should be fine. I included the screenshot of the website because I know what I'm doing in this case, and I thought that would help people understand it. But absolutely, do not tell people to click links, don't click those links. That's very bad to do. So taking screenshot is safe, uploading that screenshot should be safe.

I'm unsure how you would forward an SMS to our service. We certainly don't have that capability right now. We are thinking about how to do that for e-mail, so that if you get a phishing e-mail or a spam e-mail, you can forward that to NetBeacon itself, and it will sort of hold it for you so that you can log in and then say, "Yes, this. Send it to the registrar." The difficulty with that is similar, though, to getting the e-mail headers, which is you can't just forward an e-mail. You need to forward an e-mail as an attachment so that it's complete. And that's not straightforward, again, in every single different e-mail UI that's provided. I see a hand from Amrita.

AMRITA CHOUDHURY: Thanks, Graeme. At least it's a great initiative, I would say, to start with because this is an issue. There were two questions. One is to log into the system, you need an e-mail ID. But in many developing countries, which is mobile-oriented, people do not have e-mail ID. So is it something you're looking at? A follow-up question is, which you alluded to also, is not most end users can differentiate between the various kinds of abuse? At least if I look around also, most people will not, even educated people. So is there a way you all are looking at addressing these things? Because people look at it as an abuse but they can't differentiate what is what, and most people don't like to read. Thanks.

GRAEME BUNTON: Thanks, Amrita. I will speak very plainly to this. My sense is that there's always going to be some technical barrier to reporting abuse. I would love to lower that as far down as possible. But ultimately, it's a technical harm that requires some sort of experience or knowledge to be able to report, I think. I think. Boy, I would love to get to the place where we can just take any sort of domain name from any sort of source and say, "Yes, that's a phish," or "Oh, this is spam." But I think we're quite a ways away from that. So there's just a lot of work to be done there. For the meantime, I think there's a bit of a barrier to entry, and that's unfortunate. I would love to get to the place where it's not there.

RE phone number—and I saw that, I think it was a question in the chat from Naveed—I hadn't considered that. I'd have to think about how we would do that some more. For the moment, it will still be e-mail, but I'll

take that on board as a feature request. I've got a fun document with lots of those in the list of things to get done. So I'll take that under consideration.

JONATHAN ZUCK: Siva, go ahead.

SIVASUBRAMANIAN MUTHUSAMY: Even for experienced users with some technical background, the classification and the phishing, spam, and other malware categories, it's a limiting factor. Because I experienced an incident which cannot be technically be classified as phishing, though it had elements of phishing abuse, and there was an element of malware in the design aspects. So it's very difficult to choose one of the four categories. Is it possible that you can also give another category which is useful for people who can't classify that or something that does not fit in? And the determination about the technical classification can be made by a DNS abuse rather than by the user. Thank you.

GRAEME BUNTON: Thanks, Siva. I don't think it's quite so straightforward to identify, but we'll certainly look at that. But as a whole, this is relatively new. We've been operating this since June last year, and increasing usage of it over time. We're still learning lots about the ways that the interface needs to change and what needs to get better in identifying these harms. All of that, I think, will come over time and continue to get better. So think of this always as a work in progress. I'm never opposed to more feedback

and input on making it better. It's a free service we operate as a public good. We certainly don't think it's perfect. So all of this feedback is valuable for me to help make it better.

I saw a question in the chat I think about "Do people get informed about what happened?" No is the short answer. Right now we're looking at building a monitoring engine so that we can track what happens to a domain name after it's submitted. Did the name servers change? Did the domain status change? Did the A record change, the content on the reported page change? Things like that to give us some sense if that harm has been mitigated or not. Because it could have been mitigated at lots of layers, it could have been the host, it could have been the registrar, it could have been the bad guy finished their bad thing and turned it off. So we'll begin to track that. But it is important to note that the right answer for mitigating abuse may not be to turn it off. It could be that it was a compromised website, and so there's a small business or a charity operating at that domain name and turning it off at the DNS would be inappropriate. We always need to keep in mind that, "Hey, this domain didn't disappear," that might not be the right thing to do. I see a hand from Marita.

MARITA MOLL: Hi, Graeme. Thank you. I think this is a great initiative. As I said in the chat, I don't think it's that complicated. But what's needed is education here. For years now, people have felt they've had absolutely no agency at all with respect to dealing with what's coming into their inboxes. They've just been helplessly receiving all of this stuff and have no place to report it. It's a great start to give people a chance to say, "Yes, we can

do something and this is a way of doing it." We have to, through all of our various channels, find ways of educating people about the potential impact of this kind of a service where things can be reported and dealt with or at least you feel like you've done something to make it better. What can I say?

GRAEME BUNTON: Thanks, Marita. Yeah, sure. Working on it. I saw a question in the chat from Dave, asking about the pasting of the full URL. Absolutely. Full URL is just fine. It doesn't need to just be the domain name. In fact, generally, the full URL is more helpful because it's rare that abuse is right at the root of a domain name.

I also saw someone asking about—my apologies, I can't see it as I scroll through quickly—the statistics on how much is going through. We have that, and so we know the throughput and it's measured in the thousands a month right now, which is great, really good usage. But we won't publish stats out of NetBeacon, primarily because we don't want it to be any sort of disincentive for registrars or registries to participate and integrate with NetBeacon. They get a benefit from doing so, and increasing that linkage removes friction from the ecosystem.

The DNS Abuse Institute has a separate project called DNSAI Compass, which is worth talking about briefly, the measured DNS across the entire ecosystem—ccTLDs, gTLDs, registrars. So it puts out aggregated abuse reports right now. We've been doing that since about September. You can find that at—just google DNSAI Compass. Boy, I hope it comes up with that search. I can post the link in that chat in a sec. We will begin

publishing more detailed reports about abuse across the ecosystem, including naming registrars and TLDs that have low observed rates of abuse, and those with high observed rates of abuse. So that's going to give us really rigorous, transparent, academic almost, insight into abuse across the ecosystem, and that's going to be a better place for really trying to understand the scale and scope of abuse. I see a hand from Olivier.

OLIVIER CRÉPIN-LEBLOND: Thank you very much, Graeme. Two quick questions. The first one is with regards to the language. Are you planning to extend the service to other languages as well? So at the moment it's in English, but other UN languages or whatever. So that's one. And the other thing is how would you treat requests that pertain to domains that are under country code top-level domains?

GRAEME BUNTON: Thanks, Olivier. Absolutely, we intend to translate it. We wanted to operate it for a little bit just in English first, primarily so that we can see how people use it, where there's friction in those processes that's causing people to drop out, and just learn what we needed to change before we go to the expense of translating it. But that is absolutely a thing we're going to do.

Country codes are coming. The difficulty with country codes is that most abuse goes primarily to the registrar, as the registrar is primarily responsible for DNS abuse, not the registry. But different ccTLDs offer different information in their WHOIS output, making it difficult to

automatically identify the registrar. For example, .de, the German TLD, does not provide the registrar at all. So we're unable for any .de domain to route those abuse reports in the way that we think is right. What we'll ultimately do for TLDs like that is just rest them all to the registry and say registry, "Please distribute these for us," because they own those relationships. For the others, it's okay if we can identify the registrar. Do we already have them as an ICANN accredited registrar? Can we link those two together? So that the registrar is getting all their abuse reporting in one place rather than having it come into multiple places. So it's a little bit of complication in the ccTLD world but we've got a relatively clear path to resolving it.

OLIVIER CRÉPIN-LEBLOND: So at present, Graeme, if someone submits a report under .de, as you mentioned then, what would happen? Does the system notify them and say, "Oh, not quite yet there," or does it still accept it?

GRAEME BUNTON: It should not accept it. It should error out if you try. It should say, "This TLD is not supported."

OLIVIER CRÉPIN-LEBLOND: I think that's important because a lot of our members don't know the difference between—well, the end users generally don't know the difference between ccs and gTLDs. And especially with some ccTLDs that pretend to be gTLDs for everyone. We've seen this, and it's important. Thank you.

GRAEME BUNTON: Yeah. I don't disagree. It was just like, "Can we build the thing and get it working, and getting it out there for people while we continue improving it over time?" To me, that was a pretty obvious choice for us that we can get this working for the vast—is that true? Anyway, a lot of domains or gTLDs is a great start to get us there.

I saw Naveed commented again about a fifth category for others. There's a little bit more to the—what about this generic category? So the problem with abuse reporting right now is that most registries and registrars are getting abuse reports by e-mail, and it's unstructured. Without the sort of categorization that NetBeacon provides and having a sort of blanket or bucket category is that people are just going to choose it, and we then begin to lower the quality of abuse reports. This service is really only successful insofar as we can keep the quality high. So we can do a better job in the user interface and the automatic detection of what those abuses might be to help prevent the requirement for that sort of bucket category of generic or other, because I would love to keep that as limited as possible.

I see another hand from Siva.

SIVASUBRAMANIAN MUTHUSAMY: One observation. After the report is submitted, is there a possibility that the reporter can append to that report, offer additional information? I don't see any interface in that.

Secondly, I posted a question in the chat. Why is the registrar the focal point in all matters related to abuse? Why not the registry?

GRAEME BUNTON: Good questions. You cannot edit a report because you should send the registrar all of the relevant information when you're submitting that abuse report. A real problem for registrars is people relitigating abuse report saying, "No, you don't understand. No, you don't understand. No, you don't understand," and that's painful for them. So getting it right the first time is pretty important. When an abuse report does go to a registrar, it includes your contact information. So the sender address is NetBeacon but the reply to is yours as the user so that if the registrar needs more information, they can respond to you directly. So we're not as NetBeacon in the middle of subsequent communications.

As for why the registrar over the registry, because in most circumstances, the registrar owns the relationship with the end user who bought the domain name. They have the relationship with the registrant. So for most of these types of abuses, phishing and malware, especially where they're trying to figure out if that domain name has been maliciously registered or compromised, the registrar who's closest to that problem and the one who's ultimately responsible for it. There are some circumstances like botnets, especially where the domains haven't been registered yet, but someone has unpacked a botnet and figured out what potential domains might be registered, those can go to the registry but it's primarily the registrar that's responsible for most issues of DNS abuse.

SIVASUBRAMANIAN MUTHUSAMY: I'm sorry. I remember the question that I wanted to ask. Until five or six years ago, we talked about DNS abuse in ICANN forums. The official record was that the number of abuses reported was hardly like 10 or 20 in a million, and the number of abuses reported and recorded as such may be because such a system was not developed. It may not be 10 or 100. It could even be a thousand and a million. But now after you have started this simplified the process of reporting, from your experience so far, would you agree that the number of incidents related to DNS abuse is somewhat between negligible and low?

GRAEME BUNTON: Thank you. Great question. So I spent a lot of time on that in general as part of our Compass project to measure abuse. Compass only measures malware and phishing because those have the best data sources. All projects to measure abuse are usually consuming block lists and aggregating them and duplicating to get some sense. ICANN has a project they call the Domain Activity Abuse Reporting system, DAAR, and it's totally good, it doesn't provide a lot of details in their output, but it's technically a solid piece of work. They estimate by consuming quite a few quality sources more than we do for that Compass project. They come to a figure of about 650,000 domain names are identified as abusive in any given month, and that would be including spam, including malware, phishing, and botnets. So 650,000 on I think a total number of domain names that they see are included in the project of 260 million. So it is across the entire DNS, a tiny fraction of domain names that are actually abusive. And even if those lists are small, like

they are for sure underreporting, there's definitely abuse of names that are missing from list. Even if it's twice as many, it's still going to be less than one in 260 names that is abusive. So it's a pretty low number. Marita?

MARITA MOLL: Thanks for taking my question. Because you're Canadian and so am I, I'm going to ask what you're doing with .ca. If I reported .ca domain name, what's going to happen?

GRAEME BUNTON: There's no abuse in Canada. Unfortunately, .ca will still air to. We're working to figure that out. And so I'll be in touch. I have been in touch. We're still working together with the .ca folks to see if there's a way that we can interact. But it's probably mostly a problem on my end of getting the dev resources to build that connectivity. And if I'm being very plain, our prioritization for ccTLDs is we're the biggest with the most abuse because that's going to give us the most bang for our buck in terms of priorities, where we're able to report more to registrars and get more bad domain names taken down. CA is pretty big. They've got 3 million names in the zone, that's more than respectable, but they don't generally have quite a bit of abuse. It is quite a good zone. And so it's probably not super high on our list, although personally I would like to get it there.

JONATHAN ZUCK: All right. Well, thank you so much, Graeme, for taking the time from your place. I work quasi vacation in Barcelona. At least your family's having a good time.

MICHAEL PALAGE: Jonathan, I got one quick question to Graeme, if I could.

JONATHAN ZUCK: Sorry. Yeah, we closed the queue. But okay, go ahead. Just quick, Michael.

MICHAEL PALAGE: Graeme, just real quick. With regard to the reporting metrics that PIR puts out, I noticed that you break down the takedowns, court order from law enforcement, botnets, stuff like that. Would there be any other situations? I know you have a catch-all category called other limited content. So I guess my question to you is, are all of the times that PIR takes down a domain name reported in that, or could there be certain takedowns that are not reported?

GRAEME BUNTON: Thanks, Michael. I honestly have no idea. And part of that is because the Institute and myself work entirely externally, and so that we're focused on abuse across the entire DNS ecosystem. I don't look at any internal PIR abuse, I don't operate their abuse system, I don't look at their abuse reports, I don't manage their abuse reporting. Some .ORG abuse goes through NetBeacon and we measure it as part of our Compass, but the

actual .ORG and other PIR abuse reporting stuff all goes through Brian Cimbolic, and I don't see it. Part of that is important for the credibility of the Institute. We need to be able to work across the entire ecosystem. And to be that integrated with PIR would generate conflicts of interest, and so I just don't see it.

MICHAEL PALAGE: I appreciate that. I will follow up with Brian on that because I do think that is a best practice. I know Identity Digital has a similar detailed breakdown. Anyway, thank you.

GRAEME BUNTON: Thank you.

JONATHAN ZUCK: All right, thanks, Graeme.

GRAEME BUNTON: I'll put my e-mail address in the chat. If people want to reach out, they have more questions, they want another demo, if discovered a problem or have feedback, boy, super open to getting that. I'll be in touch again relatively soon, Jonathan, as we launched this Compass project with the detailed reporting because I think you'll find that very interesting as well.

JONATHAN ZUCK: We're excited. Thanks a lot, Graeme.

GRAEME BUNTON: Thank you. Thank you for having me. I really appreciate it.

JONATHAN ZUCK: All right. Back to you, Olivier.

OLIVIER CRÉPIN-LEBLOND: Thank you very much, Jonathan, and thank you, Graeme, for presenting this tool to us. It's always good to have you come to our calls and explain what the DNS Abuse Institute is doing. Because whilst a lot of people are talking a lot, not many organizations actually do something about DNS abuse. Well, many do but I guess the DNS Abuse Institute is at the forefront of this and has done a lot, especially for end users because this tool is likely to be very helpful indeed.

We have to move on. Our next agenda item is our workgroups and small team updates. Now, this week, there was an update on the mailing list from the Transfer Policy Review Policy Development Process. So the first update we have for 15 minutes is the one on the Expedited Policy Development Process on the Internationalized Domain Names. And for this, there is a presentation also there. I believe it's Satish Babu, Hadia Elminiawi, Abdulkarim Ayopo Oloyede who are on the call. Would it be Justine? I'm not sure who will be presenting this.

SATISH BABU: Yeah, Justine will be doing it, Olivier.

OLIVIER CRÉPIN-LEBLOND: Justine? Okay, fine. Over to Justine Chew then.

JUSTINE CHEW: Thanks, Olivier. Okay. So do I have 15 minutes or is it less?

OLIVIER CRÉPIN-LEBLOND: You have 15 minutes, yes.

JUSTINE CHEW: Okay, cool. So I'm going to be judicious today because we couldn't get through all the things that we needed to get through two weeks ago. So I plan to be judicious and just run through the main slides. Okay. If there are questions, please put them in the chat, and I'm going to ask Satish to monitor the chat and answer as best as we can. If there's any questions that cannot be answered today, then we will pick it up from the recording and get back to you offline. Okay. So I'm going to try and cover these two topics today, which is the string similarity review and what happens with string contention. Okay. I probably won't get to the other sections that we missed out on two weeks ago, but we'll try to manage that somehow maybe next week or something, I don't know. So we have a bunch of resources in the back side of the slide deck so you can have a look at those in your own time. Moving on to the next slide.

We have covered quite a fair bit of the Initial Report. The ones that you see highlighted in yellow have been more or less touched upon or completed. 4.3 is just highlighted the number because that's the one

that we tried to get through two weeks ago and we couldn't finish it. So it's still up in the air kind of thing. But today, I'm going to touch on string similarity review and string contention, as I said earlier. So moving on to the next slide.

So we've already spoken about the fact that we're using the Root Zone Label Generation Rule as the sole source to determine a variant label set. So just to recap what the variant label set is. So if you look at this diagram, which I think everyone should be familiar with it by now, we've used it quite a number of times, to run through the legend on the left, the one that's highlighted in yellow is what we call the primary or the source label. So that's the string that we would put into the RZ-LGR tool, and that spits out the variant label set. So the variant label set is the whole thing, the whole table, border in blue. So that's called the variant label set. And within the label variant set, you would have labels that are designated with different disposition values, either allocatable or blocked. Allocatable means that it's open for application and delegation. Blocked means is not open for the application or delegation. And I'm not going to go into the process of how you would challenge any of these things. That's a separate matter. So, just to recap, so you now know what is the variant label set, you know what is the primary label or the source label, you know what is allocatable and what is blocked. Again, the variant label set is dependent on the source label or the primary level. So primary label always determines a set, which is why you need to determine the primary label first before you get set. Moving on.

So, the issue with string similarity. Next slide, please. String similarity review in the 2012 round, we used what is called a visual test. So, what

happens is there is a String Similarity Review Panel. There's a panel, there's an evaluation process. What they do is they take two strings and they look at it visually to see whether it's confusingly similar. That's the term that they use. So confusingly similar enough that it would cause a user to mistake one for the other. That is the basis of that test. What they do is they do a comparison across the same script, strings with the same script, or labels with the same script, or even cross script in some cases, especially when the strings are visually confusable. So the typical example with a cross script is the triple A. You see a triple A in Latin, it looks very similar to something called the triple A in Cyrillic, but they're actually different. The string similarity review takes place prior to objection. So the idea is anything that is not caught by a string similarity review, there is still opportunity to backstop it using objection process. So that is the 2012 round.

Okay. So what we are grappling here is what happens when we start introducing variants and what is the role that variants play in string similarity review. The EPDP team had a small team. They created a small team to look into this specific issue. The small team actually took quite a bit of time to come to some kind of recommendation. I do remember having gone through this presentation in October 2022 so I'm not going to go into the specifics again. I'm hoping that people will pull it out and have a look. So I'm going to try and keep it at a very low level conceptual explanation.

So when you when we talk about Level 1, Level 2, Level 3, this diagram shows you what we mean by Level 1, Level 2, Level 3. Level 1, basically, it's just primary and allocatable. That would feature in a comparison to string similarity. Level 1 is with primary plus applied-for variant. Level 2

would be primary plus all allocatable variants, whether they're applied-for or not. And Level 3 would be everything. So including blocked variants as well. So the point being that as you introduce more levels, the comparison becomes complicated because you have increased number of labels to compare. And you're talking about permutations. Because when you compare one, say on your left hand, you have one label, and in your right hand, you actually have 10 labels, you're comparing something 10 times. So imagine if your left hand you had 100 and your right hand had 100. It becomes just basically unmanageable at some point. So in that context, we needed to come up with a adaptation of the string similarity review from 2012 round to balance between minimizing complications in terms of the number of permutations or the number of comparisons that one has do against possible harm that we want to mitigate against.

So if we go on to the next slide, this is an example. Again, this is not new. I think I showed it to you before. I'm just trying to get across why is it that we settle on what we call the hybrid model. The reason why it's a hybrid model is because we do a cross section of the three levels that I mentioned before with the only exception that we do not request for a comparison of blocked against blocked variants.

So if you go back to the 2012 round, there was no variant. Variant wasn't allowed. Variant still isn't allowed until we finish this EPDP with the policy. So with 2012 round, we were only comparing A1 and B1. So you just narrow it down to just two labels. You compare A1 and B1, that's it. If they found it to be visually confusable, then it depends. If say one of it is existing TLD, then the other one won't be allowed to proceed. If both are non-delegated yet and applied-for then they will go

into contention set. As simple as that. But when you start introducing variant labels, then you're increasing the number of possible comparisons.

So in this situation, we have come up with, as I said, what we call the hybrid model, and there's a reason for this. So what happens with the hybrid model, and in this case of the two top labels, A1 and B1, so that still gets compared to each other per last round, previous round. But now we also have to compare B1 against A2 and A3. So we're starting to introduce the variant labels. So the green box is allocatable, and the pink boxes are blocked variants. Okay.

Then we go through the motions again. So you have the root 1 where it's A1 against B1, then you have B1 against A2 and A3, that's root 3. And then B1 against A4 to A24, which is root 5. Okay. So that's one side of the story. The other side of the story, we have to do a cross comparison as well. So we do A1 against B2 to B23. And then we do A2 and A3 against B2 and B23. So the only one that we don't do is the two pink boxes. So you notice there's no line that's connecting them directly.

So we found that if we don't introduce this kind of level of comparison, then we won't catch the yellow outputs 2, 4, 4, 5 that you see on the right-hand side in the middle of the screen. By introducing this process, we're able to catch combinations of confusingly similar labels. So if we had not introduced variant labels into the comparison, then these would have been omitted and would have been let through. The upshot of it is that because we have found confusingly similar labels, then we can kind of estimate what would happen, whether they are allowed to

proceed somehow, or they are not allowed to proceed, or they go into a contention set. The idea is that if things are found to be confusingly similar, then it's a case of it would be silly to allow both of them to be delegated. That because they invariably will introduce confusion to the end user. So that's why we introduced the hybrid model because it helps to capture more combinations of confusingly similar labels, which would have otherwise been undetected if we didn't introduce them into the [inaudible]. Okay, so that's a simple explanation of the hybrid model and how it works. Moving on to the next slide.

As I said before, we are trying to mitigate harm in introducing the hybrid model. The risk and harm that we identified were two. One was denial of service. That's not DDoS. It's this lack of connection or what we'd like to say collectively is a 404 error. Now, that that particular risk doesn't necessarily cause any harm because if you get a 404 error, then that's the end of your experience, you're not sent to another place where somebody can do something to you.

But the other risk that was identified is misconnection risk. Now, misconnection risk, this diagram is, as I understand, quite complicated, but you can study it in your own time. Basically, it tries to show this. So someone looks at a particular label, the green label, and thinks that it is this. But actually, they thought incorrectly and ended up using a different label. So they end up going to a different place, different to what they thought they would get to. So there is this element of confusion and element of being misdirected to somewhere else. So that's what we call misconnection.

So anytime there's a possibility of this, there's always going to be a risk of abuse. Because you're being inadvertently directed to somewhere that you didn't want to go or you thought was not what it was to begin with. And you you're thinking because it's the label or the same string, but it's actually not. So it would cause, as screen says, possibly confusion, frustration, credential compromise, accidental exposure, and the worst case scenario, it could be maliciously leveraged so DNS abuse could happen. So those are the things that we are trying to mitigate against by way of introducing variants to the string similarity review process to try and capture as many, as reasonably as possible, to try and capture many combinations of confusingly similar strings so that they don't get delegated separately or they don't get delegated at all. So moving on to the next slide.

I think I've kind of explained this in a bit, 4.1. So that is the effect of the adaptation of the string similarity review mechanism or process from the 2012 round, which is basically to modify it to include the hybrid model. As I said before, that requires that a comparison be done against all level of strings against each other except for blocked against blocked. I said before, the purpose of it is to actually mitigate potential risks from denial service. Actually, more misconnection than denial in this, but obviously, both risks are present. Detects more combinations of visually confusable labels, I think we mentioned that. We don't request for blocked against blocked comparison. It's to avoid unnecessary complexity. I believe, as I said before, I think we presented this aspect of it before in 2022 and we kind of received report or at least we didn't receive any objections to it at that point in time.

Now, the edit point from the edit recommendations to this model now—I think it's been widely accepted—is that we have a couple of exceptions, which is 4.2 and 4.3, and they are related. 4.2 talks about the possibility of allowing the String Similarity Review Panel—that's what SSRP stands for—to decide what other variants can be omitted from the string similarity review. The omission will be based on certain guidelines and criteria that will be developed during implementation and based on research study, blah, blah, blah. But essentially, the criteria that we looked at is that it has to be on the basis of manifesting low level of confusability. There's a bunch of words there. What does that mean exactly?

Very simply explained, if you compare a Latin string or Latin label against a Chinese label, because the scripts are generally very, very different, because one is ABC, the other one, sometimes it's ideograph but it's more like a picture form kind of thing, so there is very low possibility that you will come up with a combination of Latin script and a Chinese script that would be confusingly similar. That is what we mean by giving the panel some flexibility in determining how they will omit otherwise unnecessarily redundant comparison. Because, in the example that I've given, there is unlikely going to be the possibility that Latin label would be confusingly similar to a Chinese label. So there are combinations of scripts that could fall under this kind of manifesting low level of confusability. That's one exception to the hybrid model of string similarity review. As I said before, Recommendation 4.3, this guideline is going to be developed during implementation.

4.4 basically means that it comes back to the integrity of the set principle that we have explained before, meaning to say that once a

variant label set has been identified using the primary label, that set has to be kept together at all times preserving the integrity of that set. So whatever happens to one of the labels in that set in terms of string similarity would affect the whole set. In essence, if an applied-for variant label set is compared to an existing TLD and is found to be confusingly similar some way, then the applied-for set will be ineligible to proceed because it's confusingly similar with an existing TLD. If the two groups or the two variant label sets that are found to be confusingly similar somehow are non-delegated yet but are both applied for then they will go into contention set. That's basically what 4.4 comes to. Moving on to the last slide. That would be—

OLIVIER CRÉPIN-LEBLOND: Justine, I'm going to have to really ask you to be very quick for the last slide because you're over time, unfortunately.

JUSTINE CHEW: Okay. I think string contention is quite easy. I think people can just read off 6.1 and 6.2, and make sense out of it. So I'll stop there. Thanks.

OLIVIER CRÉPIN-LEBLOND: Okay. Thank you very much for this. I do realize 10-15 minutes is very little time for you. But unfortunately, we have little time today. We have other things to cover, including the policy comment updates. Thanks for this update, Justine. We'll hear more from your group and you next week. Policy comment updates now with Hadia Elminiawi. Chantelle isn't with us today so it's going to have to be Hadia. Over to you.

YEŞIM SAĞLAM: Olivier, sorry, if I may. We also had an update for the Applicant Support GNSO Guidance Process.

OLIVIER CRÉPIN-LEBLOND: Sorry. I completely zapped this. Goodness gracious. Very bad of me. Okay. Apologies then. Let's go for the Applicant Support GNSO Guidance Process. Sorry, Maureen and Sarah. I'm a little stressed because of the time constraints.

SARAH KIDEN: That's okay, Olivier. We have a brief update. Maureen cannot make it today. Quickly, if you remember in November 2022, we started our work with the GGP and we had six tasks to complete. We've completed tasks one to five and have slowly made our way to task number six, finally. You can look at the slides later on. But next slide, please.

Basically, task six is to recommend a methodology for allocating financial support where there is inadequate funding for all qualified applicants. Essentially, we agreed that if a minimum of 10 applicants are supported through their program, we will consider that a success. But what happens in the event that the program receives 15 or even 20 qualified and deserving applicants? So we inquired about the budget and it's really not clear but we've been informed that the idea is to work with the $2 million amount which is the same amount from the 2012 round. And if we work with that amount, it translates to about 10 to 15 applications so we need to think about extra applications.

We spent our meeting on Monday discussing different options and we came up with two real decisions that we want to take further. Our homework is on the next slide.

That's the homework. We have two options. The first option is to wait until we have received all applications and they have been evaluated, and then the level of support will be determined based on that. What this means is that depending on how many qualified applicants are selected, the amount is divided equally amongst them unless an applicant specifically states that they don't require full support, in which case, they receive the support they have requested, and then the remaining amount is divided between the remaining applicants.

Then option two is to do what they're calling first in, first out continuous process. You open the application round, let people apply on a rolling basis, then they are informed if they qualify. But instead of telling them how much support they will receive, they're given a range of "We will support you up to this and it should be 50% to 75%." Then when the window is closed, the exact amount of support they will receive is communicated to them.

For both options, we need to think about when the window will close. If you recall, the Applicant Support process will start 18 months before the process for the next round. We were informed that there were discussions within the SubPro about closing the process four months before the next round, but there was feedback that some people may find out about the program in those four months. Then there was another suggestion to make it two months before the close of the next

round, but then people felt that two months is also too close to the next round. So we haven't really thought about this.

There was a discussion about using funds from auction proceeds but there's nothing concrete yet. This is homework for us to think about. Though I feel that may be in the best position to discuss next week because there's homework for staff also to look at the pros and cons for both options one and two. I think that's it for me. Thank you.

OLIVIER CRÉPIN-LEBLOND: Thanks for this very efficient update, Sarah. We have a hands up for Michael Palage.

MICHAEL PALAGE: Thank you. Sarah, thank you for the presentation. Given the limited options, I was wondering, has there been a consideration of a third option? One of the papers that I had wrote for the BRG in advance of Cancún talked about what they—I think the proposal was of an ala carte menu. And the idea there was instead of sitting there and charging the applicants a quarter of a million dollar fee at the time of application, would it be easier to break the fees down into different phases? Then that way, it would actually lower the overall cost and potentially be more inclusive. I just think with the two options that you have or the two options that have been listed as the homework assignments really limits the options of maximizing participants. I was just wondering, do you see the option to potentially raise other options for considerations or do you think they're locked in on only one of two options?

SARAH KIDEN: Thank you, Michael. If you could share the link for the paper that you wrote, it would be very helpful. We actually discussed lots of options and the working group came up with these two. Every time we would propose something, staff would say, "This is not feasible. This would not be acceptable." So there are many scenarios around, for example, reduction in application fees. We discussed fee waivers. Really, many options that we discussed. But these were the two viable options. But I will take this back to the working group and ask if there are other options, basically.

MICHAEL PALAGE: All right. I will do that, Sarah. If you could just maybe pop your e-mail in the chat. Thank you.

SARAH KIDEN: Thank you.

OLIVIER CRÉPIN-LEBLOND: Thank you very much. I'm not seeing any other hands up. Sarah, because you sent on the mailing lists the fact that you were going to be focusing on other matters, is that the last call that you're going to be part of or will you be appearing again in future calls?

SARAH KIDEN: I'm here until end of month. So I'm still here.

OLIVIER CRÉPIN-LEBLOND: Okay. So we're not going to sing you a farewell song then until then. Thanks very much for the update. Very interesting, very helpful, and thanks for all the work. Now, quickly, policy comment update with Hadia Elminiawi.

HADIA ELMINIAWI: Hi. Quickly, recently ratified by the ALAC, we have none. Open public comments, we have the PTI IANA Governance Proposal and that will be handled by the OFB Working Group. If you're interested in participating in this, please join the OFB Working Group call.

Then under review also is the Internet Service Providers and Connectivity Providers Constituency Charter Amendment, and that closes on the 26th of June and also is addressed through the OFB Working Group. Addressed through the Consolidated Policy Working group is the Phase 1 Initial Report on the Internationalized Domain Names EPDP. And the EPDP Team will be presenting the statement end of this month, on the last Wednesday of this month. Meanwhile, we are following up with some weekly presentations on the report.

Then we have the Bylaws amendment and documents to implement the NomCom 2 review, and that is also handled by the OFB. Judith Hellerstein and Yrjo Lansipuro are drafting this. Again, it's with the OFB.

Then we have the proposed renewals of the Registry Agreement for .NET. Michael Palage and Bill are going to present on this today. I don't

know if I should give them the floor now. Michael and Bill, are you ready?

MICHAEL PALAGE: Yes.

HADIA ELMINIAWI: Okay, great. Thank you. Please go ahead.

MICHAEL PALAGE: Real quick. Jonathan, I was just catching up on e-mails. Do you want to load the draft or did you provide an alternate wording? I did not catch up on all of the stuff.

JONATHAN ZUCK: Sorry. By way of explanation, Michael prepared a draft that he circulated to the small team yesterday, I guess. Maybe it was the day before. It seemed not to be in the straightforward layman's language that we tend to use for our comments. So I just took a shot at using it as the basis to do a clean draft, basically, of the points you wanted to make. I don't even know if Michael has read that one yet.

MICHAEL PALAGE: I have not, Jonathan. As I said, I was a little preoccupied.

JONATHAN ZUCK: The net of this, unfortunately, we have two different drafts that we can show you. I think it might make sense for us to resolve this in the small team and put one of them into a Google Doc for people to comment on. In one case, it's a complex draft that needs simplification. And the other case, it's a simple draft that may need some clarification. It's coming at it from two sides, how to address this. So I took an attempt to put it in our voice, if that makes sense. I don't know the best way to have this conversation now.

But if you recall, from the last meeting, Bill and Mike made a presentation about the points they thought we should make and boil that down to three main points and recommendations as part of this public comment. And it could very well be that this will migrate beyond the public comment process into advice as well. But for now, it's just a public comment that's due before too long and we want to give the ALAC a chance to read it. I don't know, Chantelle, if you've got the ability to bring up those two drafts. I'll bring up my draft, I guess.

MICHAEL PALAGE: I don't think Chantelle is on today.

JONATHAN ZUCK: She's explained that. Can you make me a co-host? I guess what I could do is bring it up and just show you the recommendations we had.

HEIDI ULLRICH: Jonathan, we're ready to display. We're ready to display it for you, if you'd like.

JONATHAN ZUCK: Okay. You have my draft?

HEIDI ULLRICH: We do.

JONATHAN ZUCK: Okay. I wanted to just focus you in on the recommendations that we agreed on. First, they're here in bold, "ICANN in conjunction with the contracted party should be for the development of a..." This looks like it's far ahead. This is not the top of it. Okay. The first issue that is addressed is the one that was raised by Evan Liebowitz when this first came up, which is the language about—verified having the ability to take down sites as a result of a court order from a government demand. What's percolated up from those conversations is the idea that we might want to document this more and disclose some of these court orders that occur. That's not currently language that's in any agreement. But the recommendation was, in conjunction with the contracted parties, ICANN should explore the development of a disclosure framework through court and government ordered domain takedown so that we become more aware of what types of domains are being taken down, what the frequency of such things are, etc. I think that's where we settled. Please correct me, Michael and Bill. In talking to the small team, I think that's where we settled on it. Because right now, it's not something that exists. That's something we can borrow from the base agreement or anything like that.

MICHAEL PALAGE: Sure, Jonathan. Do you want me to speak to that or wait?

JONATHAN ZUCK: Yeah. Go ahead. I'm just trying to go quick. I don't even mean to steal your thunder here, Mike. I was trying to find a way to get this through this as fast as possible.

MICHAEL PALAGE: No problem. And to that point there, Jonathan, based upon the comments of Zak and Evan, that was something that we felt was important to add, since most of the majority of the public comments focused on that. The one thing, what I tried to do in my more detailed discussion or a more detailed comment was to actually focus in on the letter of intent, because the letter of intent which has been amended to include both .NET and .COM talks about the obligations to look at best practices. That was one of the reasons I had asked Graeme about the reporting of PIR. I thought that that could be an example. I guess that's where the original draft went into a little more detail. There you go.

JONATHAN ZUCK: We can certainly add that. The approach that I took in this draft was to say that it's two parts. One is to look at something that was new in the documents, which is what they called on us to do, which was this language. Then the other thing we're doing is not looking at new language but taking this opportunity of a renewal, look at old language and compare it to the base agreement to determine some of the

differences in language that might be in one or the other that the other might benefit from.

The next recommendation here was that ICANN should prepare a detailed comparison of the .NET and base Registry Agreements with an eye towards the global public interest. The idea being that Mike did an amazing job that I think even he would consider cursory in the short time he was given to look at what some of the discrepancies were and things that we might want to focus on in the future. We're asking ICANN to do a more thorough comparison. Calling it a red line may or may not make sense because the documents are different documents, they're not versions of the same document. But doing a comparison and focusing on a global public interest. So that was the second recommendation. If you scroll down.

MICHAEL PALAGE: Jonathan, just on—Okay, I'm sorry. I'm just reading your document real time, though. Sorry. With regard to the red line, I do think the fact that the ICANN Board has specifically said migrating to the baseline Registry Agreement is in the global public interest, I think that is something we want to point out and don't want that to get lost in the pithiness of this document.

JONATHAN ZUCK: Yes. We probably need to take this conversation back and have it because of the limited time now. The IRP came back with different recommendations about that. So we should resolve that. I guess I'm trying to just get to the recommendations that we're making on this

call. But then we'll put something out on the list that says, "Hey, look at this document and comment on it once we've resolved whatever the deficiencies of this are." I'm sure there's plenty of them, Mike. I don't mean to imply otherwise.

The other thing that came up is this notion of language that appeared in the base Registry Agreement where the capitalization of security and stability in the context of consensus policy was lowercase, whereas it's been uppercase in the .NET agreement for many years. What we want to do is ask ICANN to explain why security and stability are not capitalized in the base RA. If it's intentional, then ICANN should explore making this update to the corresponding language with the .NET RA.

From our analysis, it feels like a significant difference between the two agreements that was introduced in the base RA. If it wasn't accidental, if it was intentional, then I think we would follow up with a request to make the same change in the .NET Registry Agreement, because it basically, at least, implies a broader mandate for the community with respect to consensus policy.

Then finally, there was this discussion about economic research. The fact that the base Registry Agreement contains language that requires the registries to participate reasonably in economic research and there doesn't appear to be a corresponding clause in the .NET Registry Agreement so I think we would request that and we've repeated our request. That as part of an upcoming five-year strategic plan, ICANN Org needs to undertake a comprehensive economic analysis of the domain name marketplace. Those are our big recommendations that hopefully are reflective of the discussion last week.

We will take this offline now that we have Mike back again to figure out where this document left things out and is deficient and we'll post it for folks to comment on. We'll have a fairly short turnaround time for that because we have to put it in front of the ALAC as well and give them time to get through it before they vote on it inside the public comment period. Look for an e-mail on this in the next day, basically, on the Listserv so that you can take a look at this document yourselves and make sure you understand it and any questions answered or modifications requested. Anything else you want to add quickly to that, Mike or Bill, since we're over time?

MICHALE PALAGE: We're good.

JONATHAN ZUCK: Okay. We have our homework and soon you will have yours very, very shortly to go through and make your questions and comments on the draft. Thanks, everyone. Back to you, Olivier.

OLIVIER CRÉPIN-LEBLOND: Thanks very much, Jonathan. Back to you and Hadia on ICANN77.

JONATHAN ZUCK: I guess I would say that's Hadia.

HADIA ELMINIAWI: Okay. Thank you. Planning for ICANN77, currently we have three policy sessions. The first one is under the title An End-User Perspective: The Next gTLD Application Window. You have a link to the suggested session. Then we have an internal session led by Cheryl about the next round of prioritization. Then we have a third Policy session. It's a cross-community session on auctions, and that's led by Jonathan. I don't know if I have time to go quickly and talk about the first Policy session.

OLIVIER CRÉPIN-LEBLOND: I think time is one of the problems we don't think we don't have. But if you can do it in 30 seconds, that would work.

JONATHAN ZUCK: I think we'd probably covered these things on the previous CPWG call. I think we're going to start having planning calls on these sessions in preps. People that have signed up to say they want to participate in the design of those sessions are going to be reached out to Doodle or whatever, and we'll get those planning discussions going. But beyond that, I don't think there's any real new news on this.

HADIA ELMINIAWI: Fair enough. Then we have the talking points. You could also click on the link. Then we have a link for the special reports. This actually includes useful links like a link for the ICANN77 schedule, the At-Large ICANN77 workspace, again, a link to the At-Large ICANN77 talking points, and the link also to the Wrap-Up session. I think that's about it to me. I give it back to Jonathan if you want to add anything.

JONATHAN ZUCK: No, I think that's good. We're over time here on your call. The plan is coming along okay. We're less than a month away so let's get excited about coming to Washington, D.C. Thanks, everyone. Thank you, Olivier.

OLIVIER CRÉPIN-LEBLOND: Thanks very much, Jonathan. And thank you, Hadia. We're now in Any Other Business. In the absence of any other business, before everyone drops off the call, let's find out when our next call will take place.

YEŞIM SAĞLAM: Thank you, Olivier. As we are rotating, our next call will be next Wednesday, on 24th of May at 19:00 UTC.

OLIVIER CRÉPIN-LEBLOND: Thank you very much. Wednesday, 24th of May, 19:00 UTC, in strict rotation. Thanks to everyone for having participated in today's call. In particular, to Graeme Bunton for his excellent explanation of what the DNS Abuse Institute has been up to. Thanks to, of course, everyone who's provided updates, and to the participants, to our interpreters and the real-time text transcription service. By the way, I fail to remind everyone, but when you close your Zoom, there's usually a survey after that. So if you've found the RTT helpful, please make it known. That's it for today. Have a very good morning, afternoon, evening or night. Hadia, sorry. I didn't ask. Did you have anything else to add?

HADIA ELMINIAWI: Nothing to add. Thank you so much.

OLIVIER CRÉPIN-LEBLOND: All right. Then it's goodbye from us.

HEIDI ULLRICH: Thank you all. Bye-bye.

YEŞIM SAĞLAM: Thank you all for joining today's call. This meeting is now adjourned. Have a great rest of the day.

[END OF TRANSCRIPTION]