JULIE BISLAND: Good morning, good afternoon, good evening, everyone. This is Julie Bisland for the recording and welcome to the DNS Abuse Mitigation PDP 1 Working Group Call taking place on Monday, the 4th of May, 2026. We received several apologies today. Please see the wiki page for the list. Alternates participating today will be Edmund Brahene, NCSG, David Hughes, IPC, Chris Disspain, RySG, and Theo Geurts, RrSG. Statements of interest must be kept up to date. Please raise your hand or speak up now if you have an update to share.

All right, I'm seeing no hands. All members, participants, and alternates will be promoted to panelists. Please watch your screen for the prompt to be promoted. Observers will remain as an attendee and will have access to view chat only. All documentation and information can be found on the wiki space. Recordings will be posted shortly after the end of the call.

Please remember to state your name before speaking. And as a reminder, participation in ICANN, including this session, is governed by the ICANN Expected Standards of Behavior, the ICANN Community Anti-Harassment Policy, and the ICANN Community Participant Code of Conduct. And with that, I will turn it back over to you, Paul. Thank you.

PAUL MCGRADY: Thank you, Julie. Hi, everybody. I am glad to be with you. I am in sunny London for the INTA annual meeting and I was saying to Julie and Terri earlier that I found a nice co-working space to take the call from, but I foolishly did not bring my converter, my power converter, so I'm just

going to be going here on battery life. I think I'm going to make it. If for some reason I don't make it all the way through the call, I hope that Nick will be able to just step right in and keep us going. So if I suddenly disappear, that's what happened.

So let's go ahead and get started. We can do a thing to PDP updates. And so we'll have staff throw that stuff on. All right, so we've got some Enhanced Working Group Participation Suggestions. The reason for this, well, we have a goal. The goal is to keep fairness and predictability, but add facilitation methods that better structure discussions during working group calls. This is born out of a concern that we have a very small fraction of the working group that is participating verbally in these calls. I'm very happy and I welcome their participation, but I do wonder whether or not we're hearing from other folks and I do wonder whether or not other folks think they can't participate and/or shouldn't or whatever

So we have cooked up some ideas and I wanted to give everybody a heads up that some of these things could start to happen. Idea number one is that to proactively invite some of the quieter or I kind of like underrepresented voices because this is a representative PDP, but some of the quieter voices and allow or request queue deferrals so that participation is not dominated by the fastest or most frequent speakers. Ask for other voices that have not spoken yet on any of the calls and to give some space, some time for reflection.

Pause the queue when valuable issues emerge, requesting targeted follow-ups before moving to the next speaker. Potentially introduce tools such as breakout groups, round robins, live polling, and report-

backs to broaden input and surface more perspectives. Anyone who's not yet spoken might be allowed to jump ahead of repeat speakers in the queue. That's within reason. Sometimes there is some back and forth that needs a quick back and forth between the people who have presented the idea and the people who are refining it.

It might rotate first speakers, even if they raise their hands, so that the first speakers don't frame every issue. And then I'm considering round-robining for major decisions or topics. So, you know, I may sit down and say, okay, every single stakeholder group is going to say something. Now you get 30 seconds and read down the list of stakeholder groups. This is not a criticism at all of the folks who do engage a lot. I think they are driving the momentum and energy of this PDP, but at the same time holding that intention with the fact that we would like to make it as easy as possible for others who maybe are hesitant to come forward and start to share their views.

I think that we're just trying to find ways to do that, because we are going to end up at a consensus call and that is not the time to say, well, you know, I never raised my hand and never spoke about these things and so now I want to share my opinions at the bitter end, because that's not a good PDP outcome. So anyways, this is just something that we are kicking around in leadership and staff and wanted to alert you guys if suddenly, you know, the queue starts looking a little bit different or those kinds of things, that's what's going on.

Let's move on to the next slide. The idea of -- oh, we have a hand from Dennis Tan. Before we move on to the next thing. Dennis, go ahead.

DENNIS TAN: Thank you, Paul. This is Dennis, Registry Stakeholder Group. Just one quick clarification question. When you refer to voice, sounds like you are equating voice to person instead of voice to stakeholder group or group constituency, or what have you. We, I mean, just for the sake of the example, Brian, Chris, and I represent the Registry Stakeholder Group. So either of us, when we speak, we speak on behalf of the stakeholder group. And we might say, you know, use our own words, but we'll speak the Registry Stakeholder Group's position.

So our lack of actively being the mic does not mean that we are not represented, if that makes sense. So if I don't have the same airtime as Brian, that does not mean that I feel underrepresented or the Registry Stakeholder Group is not being represented. And that's what I'm saying. So I'm not sure whether that exercise of calling people out to speak will enhance the conversation. Just an observation.

PAUL MCGRADY: Okay. Yeah, no, I appreciate that. I am not meaning to equate individual voices with voices. This is a representative PDP. And if we can go back to the screen, I think we tried to capture that when we looked specifically, like for example, you know, the round-robining for each stakeholder group to give a summary. Right. But yes and, yes. And if we say, "Does anybody from the registry group want to add anything to this idea before we move on?" And you guys say, "No, we said it all through Brian," I think that'll be just fine.

All right. Let us go on. So we have talked about this to make the working group move faster, including the introduction of glossary. We are starting that based upon the strawpersons developed so far. The glossary will be included in the collaboration document going forward. And we are asking working group members to review and react to comments on the document, like just directly into the document. We're not going to spend working group time going through those kinds of definitions, but we encourage you to look at it and add your comments.

The purpose is to create a shared baseline understanding for the working group across key terms. And it will help external readers of the initial and final report to navigate it when it comes out. And we think a clear common vocabulary will make it easier to understand the recommendations ultimately when ready. So, look for it. Staff is going to kick it off in the co-lab document. Look for it there. And don't be shy. Put in your comments on it, including if there is a term you would like to define this, not in there, throw it in there and put in a proposed definition. If there's an existing definition and you want to tweak it, put that in also as a comment and we'll take a look at it.

All right. We are moving on to Discussion of Strawperson Charter Question 3. We've spent quite a bit of time on this, but we think it is worth revisiting as we continue to refine it. Defining “investigation”: what constitutes a reasonable investigation by a registrar? What investigation steps are required or recommended? Are the criteria for investigation proportionate and necessary? What is the impact of this investigation on domain name registrants?

So let's jump into the strawperson. So a reasonable investigation by a registrar must consist of reviewing information reasonably available to the registrar in the normal course of its operations to determine whether associated domain names are being used for DNS Abuse per 3.18.2, I think, of the RAA. A reasonable investigation must be practical and proportionate based on the circumstances, and consistent with Section 3.18.2 of the RAA. A reasonable investigation must not require registrars to access or generate data that is not reasonably available to them, recognizing that registrar capabilities, technical systems, and available data may differ across and within registrars.

So what investigation steps are required or recommended? A registrar must, at a minimum, review at least one reliable internal data point reasonably accessible to it and may use additional technical or abuse intelligence signals where needed or available. This policy should not prescribe a fixed or exhaustive sequence of investigation steps, but should allow registrars to apply a range of indicators and pattern recognition signals appropriate to the circumstances. This policy only requires ADC within a registrar portfolio and does not require cross-registrar, inter-registrar checks and coordination. A reasonable investigation should remain distinct from enforcement, meaning registrars should assess available evidence first and determine proportionate next steps based on that assessment.

All right, I think that's it. Yep, one more. Oh, here we go. Are the criteria for investigation proportionate and necessary? Registrars must apply association criteria in a reasonable and proportionate manner, considering the risk of over-association, including scenarios where domains may be linked through reseller models, privacy proxy services,

or shared data that does not reflect common control. The need to avoid resilience on assumptions that all domain name sharing attributes are not necessarily controlled by the same actors.

The importance of ensuring that association is based on indicators that are meaningful and actionable rather than speculative. The importance of accessing data only when deemed necessary to perform the ADC. Usage of minimal information reasonably available and necessary to determine whether domain names are associated and avoid collecting or processing additional data not needed for that purpose, and the potential impact on domain name registrars.

All right. What is the impact of this investigation on domain name registrants? A reasonable investigation should be designed to minimize unnecessary impact on legitimate domain name registrants while enabling effective mitigation of DNS Abuse. The working group recognized that overly broad or speculative investigation requirements could create unnecessary impact on legitimate registrants. Insufficient investigation could fail to identify coordinated abuse campaigns. A balance is required between effective abuse mitigation and minimizing unintended consequences.

All right. I think that's it. Staff will correct me if I'm wrong, but I think that is. And we have working group comments on the charter questions, yes. Farzi, I see your hand’s up. I promise I'll get to you as soon as we get through this reading.

Working group comments in the collaboration documents on strawperson for Charter Question 3. There is some rationale for moving

some language from question three to question two and discussing how to make question two and three more coherent. Some working group members discussed whether a registrar must, at a minimum, review at least one registered name holder or account-related indicator versus check a reasonable amount of indicators until it is enough information to decide whether other domains are associated with the specific case.

We also kicked around the question, are one registrar's reseller information reasonably available to them, because resellers may hold additional data points crucial to a comprehensive ADC. Ideal scenario is for the registrar to investigate in cooperation with their resellers and other participants of their business model.

And then we looked at "at least one reliable internal data point." And some folks said that they worry that this dodges the requirement to check whether or not the threat actor has multiple customer accounts. For example, hypothetical registrars, one reliable internal data point, they choose is looking at all the domains in the customer account. They then don't take the additional step to see if 10 other accounts exist with the same payment information, name, IP login, name servers, etc. And the question is, in such a case, is the registrar compliant? All right. And is that it, staff?

JOHN EMERY: Yes.

PAUL MCGRADY: Okay, we did it. I'm sorry for being in a little co-working box with only one screen. Now let's go back up to the top of the strawperson. And the goal of this conversation is to just sort of, you know, figure out what else with this, if anything, needs to be tweaked so that we can move this into stable but not final form. Farzi, your hand was up, so go ahead.

FARZANEH BADIEI: Thank you. I can't remember my comment. So I think it was about the end user impact slide. If we could go to that slide. So when it says insufficient investigation could fail to identify coordinated abuse campaign, I'm sorry, it's domain name registrant. This is not an impact. The second bullet point is not impact on the domain name registrant, it's impact on the end user. I know that we are not talking about the end user here, but you can like put that bullet point somewhere else.

I also don't think that a balance is required between effective abuse mitigation and minimizing unintended consequences is enough; we need to come with a more granular language, because we are literally asking for profiling and looking at all the accounts and stuff like that. So if I may, we will talk to NCSG and see if we can come up with some more granular language that is not necessarily binding, but can clarify the impact.

PAUL MCGRADY: Okay, yep. I see no harm in that. I think that maybe is going more granular than this document anticipates, but let's see what you guys come up with. Theo, go ahead.

THEO GEURTS: Yeah, thanks. And this is Theo for the record. So reading the language, it looks pretty good. I don't see anything that sort of jumps out here that we need to address. What I do like to add is here, and this is just going back to the previous discussions about this charter question: will there be registrars who are going to try to wriggle out of their requirements? Sure, that will happen. Will that last? No, it won't. We got to keep in mind that we are living in a world you where we are data driven. ICANN OCTO knows exactly which domain names belong to clustering registrations. While they don't know the registrant, they do know with a very high confidence which registrars belong to X registrant.

Now they are not the only one. I mean, all these cybersecurity investigators, be it ChainPatrol, be it FishDestroy, it doesn't really matter who. They also know what these cluster registrations are. I mean we've already seen the evidence when .top was not in compliance. 200,000 domain names were taken offline in that process, and if you looked at the data sets that were made available, those were all cluster registrations. So there is no way in the end that you will go to wriggle out of these obligations as a registrar. You will be audited or you will be audited by some security forum who will provide tons of evidence to ICANN. Thanks.

PAUL MCGRADY: Thank you, Theo. And we have Gabriel next.

GABRIEL ANDREWS: Hi, can I ask that we flip back to that slide that had the feedback that you had numbers one through four? I just wanted to comment because I think three and four both were some points that we had interest in. So, if you don't mind, let's just tackle both of those. So number three is first asking about whether or not the information is reasonably available to a registrar when it's talking about when the information is in a reseller's hands and not the registrar's to start.

I think there was some margin chat from my counterpart, Naum, asking this question in the margins, and then I note that Reg constructively had provided thoughts on a potential for the reseller itself to be the one conducting the associated domain check when the business model makes that the most appropriate course of action.

I don't philosophically have any objection to that, but I don't see any treatment in the text as is for what the requirements would be when the reseller is the one making use of the reseller information noting that the registrar's still on the hook. But I have a little bit of concern with regards to the text that says a reasonable investigation must not require registrars to, you know, or something to that effect, or must consist of information reasonably available to registrars. Actually, that's it in the further text.

If the check is something that must consist of information available to the registrar, well, how does that work when it's the reseller that has that information? And so I put an offer. Reg, I don't know if you saw it in the margin, but if the Registrar Stakeholder Group, especially those that work with resellers a lot, have text that they would like to suggest that encapsulates the expectation that the reseller might sometimes do

that ADC on behalf and somehow, some way have the justification component come back up through the registrar to be able to provide to ICANN compliance, I would be very grateful if you could suggest such an edit. I'm going to pause in case there's a reaction there, but then I would also like to tackle number four next.

PAUL MCGRADY: Great. I saw Reg's hand go up, so Mary, be patient with me. We'll have Reg respond directly to Gabe, and then, Gabriel, you can keep going on four, but let's keep the flow going, if that's okay. Reg, go ahead.

REG LEVY: Thanks. This is Reg Levy from Tucows. So, a reseller will have information that a registrar does not have, but a registrar will still have information. So, regardless, the duty falls to the registrar to ensure that the obligation is fulfilled. So if the registrar is fulfilling the obligation, then they will look at the information that they have to them. If the registrar is passing through this obligation to the reseller, then the reseller will use the information that is available to them. And I think this is probably closer to what a lot of people want in the case of wholesale registrars because the reseller has better information available to them about the specific person who registered the domain. And no, I won't be sharing contractual provisions.

PAUL MCGRADY: All right. Thanks, Reg. All right. Mary, you have a comment on resellers as well, before we go back to Gabe for number four?

MARY PENN: Yes. Thank you. Mary Penn, representative from the IPC. So we've mentioned this before, but I mean, why would we set up a structure kind of allowing for a way out for a reseller model? So if a registrar chooses to have a reseller model, there's still an obligation per the RAA to pass down and flow through all of those obligations. So this kind of plays into the intra-registry check. Why would we have an out? That doesn't seem to make sense. That's not the reasonable approach. There seems to be some misalignment with the text and then what we've talked about being the objective of the PDP. So can we help right-size that?

PAUL MCGRADY: Thanks, Mary. Where specifically, if it's okay, or you don't have to do it right now, we can do it on the list or wherever, but if we take a look at the first slide, which are sort of what are going to become draft recommendations, it looks like, is there a specific text you could help us with either now or in the collaboration document where you can capture that idea?

MARY PENN: Yeah, I'm happy to put that into the collaboration document. But, I mean, it's the last sentence in that first paragraph. So, must not require to access to generate, you know, access not reasonably available to them. So how are we defining what is reasonably available? We've talked about the reseller model should not be a catch-all to, you know, allow for not checking and doing an ADC just because it's within a

reseller model. So that's the agnostic model that we've approached or we've recommended not to use as an agnostic approach just for this very reason. I guess I'm not understanding how that fits the reasonable description or if we're going to tighten that up just to kind of prevent loopholes.

PAUL MCGRADY: Gotcha. Thanks, Mary.

GABRIEL ANDREWS: Sorry, just to jump in again, but I want to expand on the text. I put it in the chat, too, just for your awareness, Paul. I pasted it with an @Reg at the start, but there's the text also that says a reasonable investigation by a registrar must consist of reviewing information reasonably available to the registrar. And so, again, here, like it seems very focused on the information available to the registrar without giving any consideration at all to those circumstances where the registrar doesn't have that information. I think that's really the key choke point that we're seeking to address.

PAUL MCGRADY: Okay. Thanks, Gabriel. I saw Reg's hand go up during Mary's intervention and a couple of others. So, Reg, if you want to speak to this issue, let's go ahead.

REG LEVY: Thanks. I don't see how this is an out for wholesale registrars. To me, this actually enriches the ability of a reseller to perform a check. So some of our customers prefer to do the checks themselves and others would not. My team would still be doing a significant number of these checks and where we see patterns across our registrar within our resellers, we're going to be seeing those patterns and taking that action anyway. But to the extent that abuse gets reported directly to a reseller, they will be able to look at, for example, payment data that we just don't have.

And so what this to me says is I don't have to go seek out payment data because I don't reasonably have it. So I don't need to check or create or receive or whatever payment data because it's not reasonable for me to do so. I look at other data points. And to the extent that my reseller is performing the obligation on my behalf, they look at the information they have reasonably available to them. In any case, I'm on the hook.

PAUL MCGRADY: Yeah. So I think we have agreement on that concept, right? That the registrar is not just saying, “Oh, we have a reseller model, I can't help you,” is not going to be good enough. I think we just need to bake that in here and then iterate the text until everybody's happy with it. So please make comments in the collaboration document and staff and leadership will also look at that on our next call. We have Theo and Volker, if you're speaking to this issue. Mary's back. We'll hear from Mary again. And then if we have exhausted the issue, we'll go back to Gabriel for his next item, but not until we're done, because we do need to get to clarity on this. Theo, go ahead.

THEO GEURTS: Yeah, thanks. So during investigations, we often reach out to resellers like, “Okay, these are our findings, but we are lacking this and this information. Could you do an additional investigation and use the information that is available to you?” That's already happening. What I can do is I can request that information. I would love to as an investigating point of view, but I can't ask for the IP addresses. I can't ask for the payment data. I mean, there's all kinds of regulations around payment data. And of course, you've got things like GDPR when you talk about IP addresses and so on and so on and so on. I think we need to stay clear from that subject or this working group is going to get stick in the mud on GDPR. Thanks.

PAUL MCGRADY: Thank you. Okay. We have Volker next and then Mary.

VOLKER GREIMANN: Yes, thank you. And I disagree strongly with terming this as an out. This is not an out. This is just recognizing business realities. We cannot be expected as a retailer registrar to have the same data available to us that a retail registrar has. We do not have account data of the registrant. We do not have billing data or fulfillment data. And thus, all we can do is, when we forward a complaint to a reseller to investigate, is ask them and have some contractual obligations in place that basically requires them to do the same that we are required to do. Putting obligations on registrars that we cannot fulfill, that are impossible to fulfill, is not a good idea because that will make the policy

unimplementable. Ultimately, we want to have something that can be enforced and will not run into a roadblock on the first stop. Thank you.

PAUL MCGRADY: Thanks, Volker. Mary, go ahead.

MARY PENN: Thanks, Mary from the IPC. So the concern that I have, and I'm still not understanding how we're getting here, limiting the scope and any sort of obligation to what is available to you. So it seems to be, I know there's been disagreement that there's not an out. It seems to be an easy out to say the information is not available. That lies with the registrar, or as a reseller.

But then further, we're limiting, and specifically in the text of this document, intra-registry checks. So there seems to be a large disconnect with where the data lies and any obligation to try to access that. So I do see that this is something we can fix with the language, but just because there's a reseller model should not be an opportunity to say, well, we don't have that information, it's not reasonably available, when it's required to be collected per the RAA. That's the disconnect that I'm having.

PAUL MCGRADY: Gotcha. Thanks, Mary. Eberhard and then Volker.

EBERHARD LISSE: We are here to write policy for the registrars, not for resellers. While I agree with Volker, it may be difficult to enforce, we shouldn't be over prescriptive in the language; it's the registrar's responsibility, they have a contract with ICANN, or the agreement; that agreement says they must do certain things, so from that follows they can have an agreement with their resellers which says they must do certain things. And we shouldn't be overly prescriptive in telling exactly what to do, but the registrar is on the hook like Reg says, and I agree with this, and they will in the end, if they're on the hook, and compliance can look at this as -- I forgot who the colleague was who mentioned the OCTO that can be done.

I have a little bit of a concern about something that wasn't mentioned. It's not enforced anymore. It's a possibility. I don't want any registrar to be forced to buy out information from third parties; especially smaller registrars will have problems on that thing. And finally, I thought we agreed previously that we won't interrupt each other.

PAUL MCGRADY: Yeah, thanks, Eberhard. I appreciate that. And yeah, we are going to do our very best to stop doing that. Volker and then Reg.

VOLKER GREIMANN: Yeah, maybe I'm stupid, but I just don't understand how we are able to do ADCs on data that we do not have. That seems like an impossibility by definition. We can only do ADCs on information that we have. We have the information that we are contractually required to. However,

there's other information as well, like billing data, like account data of the ultimate registrant. We are not required to have that by the RAA.

That's data that the resellers will have, or the resellers of resellers will have, or the resellers of resellers of resellers will have. Why would we be required to look at data we do not have? We only can look at the data that we have. And everything else suggesting that even is just making a big mess out of this.

PAUL MCGRADY: All right. Thanks, Volker. Reg, go ahead.

REG LEVY: Thank you. I'm saying essentially what Volker is saying, but with a bit of a twist. The registrar has all of the information that the RAA requires us to collect. Our resellers collect it and send it to us. So we can review all of that data. There may, however, be additional data that a registrar may have, a retail registrar or a wholesale registrar that it wants to review, that it could review, that it should review.

However, a reseller may also have a lot of that data. So a retail registrar is not required to collect payment data because that's not in the RAA, but may do so anyway, and may want to review payment data with regard to an ADC. A wholesale registrar does not have payment data, so they cannot review that, but a reseller may and should review that if the reseller has the pass-through obligation.

The reason I'm choosing payment data is because it's very specifically something that's not collected by the RAA or required to be collected by

the RAA. So what we're saying here is, and I think Theo touched on this as well, IP address information, again, may be something that a reseller or registrar collects or it may not be. So if you are even a retail registrar and do not collect IP address information, you're not going to be required to check it.

I also put into the chat, but the chat is difficult to navigate since we don't have responses available, so I'm going to put it in again. This is a link to a graphic that the Registrar Stakeholder Group put together that describes some of the data flows and some of the information that a registrar may have based on various different types of registrars. We go through retail, corporate, and wholesale.

PAUL MCGRADY: Thanks, Reg. It occurs to me that I don't know that I've ever seen in one location a exhaustive list of the kind of data that is required by ICANN for registrars to have. I think that would be useful. So we're going to make a footnote and see if we can get staff to dig that up for us or make one. We have Benjamin.

BENJAMIN AKINMOYEJE: Thank you. So I just wanted to ask that, how easy or difficult is it for the registrar to seek for additional information from the resellers if they need to, in order to be able to fulfill the ADC in any situation, rather than asking the reseller to always collect the data? Because we are of the opinion and supporting Farzi's point of view, is that we should not ask the reseller to be collecting all those data, especially the ones that violate things like GDPR and such, or privacy. So, wouldn't it be easier

for the registrar, if they need to, to just ask the reseller for such data if they think they have it?

PAUL MCGRADY: Thanks, Benjamin. That's a good question. I saw Volker's hand go up. Maybe he'll address that in addition to whatever he was intervening for. I see Carlos is popping up, so we'll get to you. Volker will wait. I want to go down the line, but I'm just prepping you. All right. We had Gabriel next, and then Mary Penn and then Volker.

GABRIEL ANDREWS: Hi. Okay, so thank you. And I wanted to address that I think the GAC isn't viewing this as an out per se, Reg, but rather a potential flaw in the structure of this policy language that needs to be addressed in order for the ADC to achieve the policy purpose that we're putting on it. And so I put in the margin again in the actual document to capture it because I don't trust our chat to really capture this very well.

But I think the fundamental question that we're asking that we just hope that the language can be revised to address this. So we don't have to answer this question right now, but like we want to see the language end up answering this question. If the reseller is the only party that's capable of doing the associated domain check based on the information that they have, payment information is a great example, how is that associated domain check obligation expected to be enforceable against the registrar? And really, that's what we're looking to have a clear answer to. Over.

PAUL MCGRADY: Thanks, Gabriel. All right, Mary, go ahead.

MARY PENN: Yeah, so I mean, I understand what's being said, and I hear the point you can't use information and be required to check information that you don't have. What my concern is, what we're saying doesn't seem to be in alignment with the text, and so I could take some time to kind of make some suggestions. But when Reg just said that there's information that they should be checking, even though it's information not required to be collected on the RAA, where is that here? It only says their obligation is minimally one data point, but it doesn't account for any information that they might have that Reg herself said should be checked, that's my point.

So there seems, yes, there's practical limit limitations, understand that, but we shouldn't apply that practical limitation across the board when it may not in fact really be a limitation. So I think we need some more flexibility and allow for that to be reflected here in the document as an obligation on the registrars.

PAUL MCGRADY: Thanks, Mary. All right. We have Volker and then Carlos.

VOLKER GREIMANN: Yes. So ultimately, you want to ask yourself, how big a gift do we want to give to abusers with this? If you are proposing that we have to reach

out for every abuse complaint that we get that is with the reseller to the reseller and ask them individually to provide certain data for that, wait for that data, then look at the data, we'll be spending weeks on every single abuse complaint, not minutes as we're now.

I think we need to be realistic in what can be achieved, what can be operationalized here. The ADC must be something that can be conducted as part of the investigation against the domain name itself. If we are spending hours, days or weeks for feedback loops that may or may not work out in the end, then I think we missed our goal here. And then we're giving a big boon to the abusers because ultimately, we'll be doing their work for them and delaying our abuse handling processes to grind them to a halt.

The ADC as we envision it and as we currently operationalize it, is something that is always based on the data that the registrar has available to it. That is something that we can immediately check on. That is something that we can incorporate in our standard processes and make sure that we find other domain names.

And some registrars may have some data available that other registrars don't. Some registrars may have more information, some may have less, but that doesn't mean that any registrar is not obliged to do ADCs. They just might do it differently. And that is what we are trying to achieve here, I think. If we try to achieve a process where every single registrar has to do basically the same check, if it's through resellers or not, then, frankly, we are absolutely heading for chaos here.

And I think we need to make absolutely sure that whatever we come up with can be operationalized, must be implementable in a way that doesn't interfere with regular abuse handling processes and does not grind these processes to a halt or delay those, take up capacity more than absolutely necessary to do an ADC.

Last week I had a session with Mark where we took down a domain and he had asked me about, and I contacted an ADC right there. And that was a domain that was with a reseller, but there was information rightly available that may help me identify a number of domain names with an account.

There's domain string similarities. There's the registration dates. There's hosting providers that a a certain number of domain names may be pointing to that when taken together indicate that there's an association there. And that doesn't require me to do a billing data or account data or registration data in the least bit. That is ADC as well.

Some may have more abilities, some have less, but I don't think that we need to be overly prescriptive on what exactly the ADC is. There should be a minimum that should be defined, but over that minimum, I think we will come into a situation where it's just not operationalizable. Thank you.

PAUL MCGRADY: Thanks, Volker. All right, we have Carlos, go ahead.

CARLOS BECERRA: Carlos, Registrar Stakeholder Group. Yeah, just kind of plus one to everything Volker said. Every registrar should have enough information to conduct a complete ADC. Any additional data would just be supplemental data and more directional than anything else. Also want to point out that just because this sets the minimums, it doesn't preclude any registrar from reaching out to the resellers or their partners to obtain more data. If the original investigation leads them to seek that additional data, we shouldn't go harvesting data. We should only go and seek data if the investigation leads us in that direction. Thank you.

PAUL MCGRADY: Thanks, Carlos. All right, Reg, go ahead.

REG LEVY: Thanks. Gabe said something that I responded to in the chat, and now I'm seeing that a reseller is the only party that can perform an ADC. That's absolutely not what I said, and that's certainly not what I meant if that was what was understood. A reseller may perform a different ADC than a registrar might perform. The registrar is still on the hook. So the registrar may still be performing an ADC, but it may be a different one, as Volker said, than what a reseller does.

So what is listed here is a reasonable requirement for the party conducting the ADC. And as Ching said, the registrar may see things across resellers that a reseller may not see. And none of that is precluded by this. In fact, what it does is it allows us the flexibility to do that without requiring that a retail registrar do something that it cannot

do and without giving the bad guys a roadmap about how to get around this.

PAUL MCGRADY: All right. Thanks, Reg. Volker, is that an old hand or are you back up? Old hand. Okay. All right. Any further thoughts on this from anyone who has not spoken or any stakeholder group that has not shared an opinion before we move on? Rod, go ahead.

ROD RASMUSSEN: Yeah, thanks. I was waiting for this whole reseller thing to kind of percolate through. One of the things I'm concerned about the way this is written, I don't think it's intentional, but I'm just concerned I'm raising, is that this could give a way out of taking into account information that is actually reported to the registrar or whoever the party is looking at this about what the association is.

So I want to make sure whatever language we create here does not preclude. In fact, I would say the opposite would, to the extent possible, require you to take into consideration information reported as part of the report that generated the initial action that would led to the trigger that you have a pattern that is being reported to you by the abuse reporter that you take a look at that.

So in other words, we don't limit things to just whatever the registrar may have on hand or the reseller from the information that they collected in the course of normal business, but they also take into

consideration information reported to them, like here is what the bad guys are doing and here's what to look for. Thanks.

PAUL MCGRADY: Thanks, Rod. And I think that's a good summary of this because it seems to me that there is a certain data set that's required by the RAA, and that's a knowable thing, right? And so we should know that. And then there is data that the registrar may have, even if it has a reseller model that was reported to the registrar, came to be known by the registrar. That's data that it has access to, it doesn't have to go and create.

And then there may be data that the reseller has that may not technically be part of the required data set under the RAA, but it might be good data in terms of being able to conduct the check. So once we identify the world of the kinds of data, then we need to set the must and the should or the may knob on each kind of data, so that everybody knows what the responsibilities are.

And again, just because a registrar doesn't have access to a kind of data that a reseller may have access to, that doesn't mean that the registrar, you know, should not ask for it in case they think there's some benefit to having it. But we can set the knob. I think this has been an incredibly good discussion because we've been fuzzy about data, like the required data set and other things the registrars have access to anecdotally versus what the reseller has in these circumstances.

And now I think that we can figure out, as Jothan mentioned in the text I saw, that there is a floor and a ceiling. We want to set a floor and

ceiling thing here that we can all work to. So I think staff and leadership have some homework to do to put together that, at least an initial understanding of the required data set under the RAA, it's a good starting point; and so you'll see that pop up in a collaboration document or in an email or something like that, and we're going to encourage people to weigh in on that. And I think that will help us get this strawperson across the finish line. I'm declaring it stable but not final, I don't think we're quite stable yet. I still think there's more work to be done here, but I thought it was a good conversation.

Now, I think Gabriel had one more point on the feedback section. And I know we're taking -- yeah, I was like, Gabriel, you're going to get me in trouble with Eberhard because you interrupt me, so don't do that. But I'll call on you. But I think we should do that. I realize we're taking a big chunk of our time. But the conversation was not circular, it wasn't pontification. It was drilling down, and I thought it was good. So thank you all for being patient with me. Okay, Gabriel, floor is yours.

GABRIEL ANDREWS: Yeah, and apologies, I thought you were calling me just then, so I didn't mean to over speak you. So with regards to the separate topic of bullet point number four, this was a comment that I had in the chat and I might be reading too much into this and so I just wanted to call that out. But the language that I was capturing and responding to mentioned that a registrar should be making use of at least one reliable internal data point. And I wondered in the hypothetical situation, like, so the question here is, is a registrar compliant in this hypothetical that I just

lay out? If a registrar chooses to look at all domains in a single customer account when they receive an abuse report.

I mean, that I would consider might be a reliable data point to identify additional abuse. But if that's where they stop and they don't check to see if they have 10 additional accounts that have the same payment information, the same name, the same IP -- well, let's say that there are, in this hypothetical example, there are 10 additional accounts that all have the exact same payment information, name, IP login, name, servers; is the registrar in that case compliant with this policy? And I don't know the answer to that based off of this text. And what are folks' thoughts on that?

PAUL MCGRADY: Thanks, Gabriel. Let's open a queue on that. All right, Theo, go ahead.

THEO GEURTS: Yeah, still trying to distill the scenario here. Speaking personally, I don't see an issue there because the way we structured our database, it's going to be very easy to do this. And within this conversation, what I usually do is I take the email address, which is always supposed to be a unique data point. And then it will pop up. I mean, if there are different resellers involved, if there are different registrants involved who use the same email address, I will get those results.

I can't speak for the other registrars because this is not a simple search function within a database. Would it be compliant or non-compliant? I think that is a question that will be shaped by the future. If I would use

the email address as the reliable data point in my scenario and I don't come up with the correct findings, or let's say the reporter has already determined that there is a cluster and my search doesn't pop up with the desired results, then I will get corrected by either the person -- the reporter goes to ICANN compliance, ICANN compliance goes like, well, there is this certain amount of evidence. Why didn't your search on that particular reliable data point that you did used -- why didn't it pop up?

Then I need to answer ICANN Compliance. Then I might find like, oh, that wasn't a reliable data point. I should have used the phone number. If I had used the phone number, then I would have found all those dodgy domain names and I would have taken action.

Because when you do these investigations, I always say they're like a 50,000 shades of grey when it comes to reports about dodgy, malicious phishing domain names. There is no real hard scenario. Sometimes you need to do a deep dive and sometimes the criminals make it very easy and most of the times you just can go with the email address and you capture like 99% of all the domain names that are clustered. Thanks.

PAUL MCGRADY: Thanks, Theo. All right, we have Farzi.

FARZANEH BADIEI: So what Gabriel is saying here is exactly why we mentioned that we were like arguing over the trigger point, like if you have one actionable report, then the registrar will have to do all sort of kind of like profiling

activities. And that's why we kept saying that maybe it will be better to contextualize and maybe leave it to the registrar or see if we can come up with general indicators that at which point which data point can be used and should be used for the ADC.

But that didn't get anywhere. But it just shows that -- so we are like -- you know, if it's like a severe harm that is being imposed and like there is a botnet attack, obviously looking at the different accounts that have received one DNS Abuse report makes sense. But in other circumstances, it might not, and it might actually put the end user and the domain name registrant at risk. And we are coming up with several examples of that.

So I think that we are comfortable with having one indicator at the moment because we could not limit the investigation and decide on the proportionality on the investigation based on the trigger and the other indicators. So, yeah, that's exactly why we are worried about the one abusive domain action. And I think that these questions should be bundled together and answered together.

PAUL MCGRADY: Thanks, Farzi. So, I mean, ultimately, all the draft recommendations will be read together. I know we've spent some time on the issue of if we should have a trigger that even though you have a report of abuse, that the trigger isn't triggered. Like I said, I don't think that really got anywhere, but I do think that out of that conversation, you've raised some important issues about contextualizing what happens, you know,

in the process. And so I do think that we need to capture that and think that through. Martina, and then Gabriel.

MARTINA BARBERO: Thank you very much, Paul. And I will be very brief. I know that Gabriel will come in after me, but just to say that I think what the GAC put in the text is this light concern that we know that most of the registrars will apply ADC and they will even go maybe the extra mile in applying it in a way that makes sense. But if we set this bar of only one reliable internal data point, is there any way for a registrar who does not want to do any ADC to just pick every time the data point that makes the least sense to check to find any connection and association?

And that would be considered as compliant in terms of the policy, because ICANN compliance would say, yes, they have checked one internal reliable data point, and that's it. But if I play this game, I mean, this is why the GAC is not super comfortable maybe with one reliable data point, if it allows to play this game. But this is where we also welcome clarification from the registrars if they feel this is something that is likely, it's possible. And I think the scenario that Gabe was pointing at was linked to that in terms of, you know, what would be considered compliance. Stopping here.

PAUL MCGRADY: Thanks, Martina. It's interesting. Yeah, unfortunately, we always have to be concerned that our language doesn't create loopholes. I think that we should figure out what we're trying to get to by saying at least one reliable internal data point. And I think it's written under the

assumption of good faith because we hang out with good faith registrars on this call and in our ICANN meetings, but we probably have to take a look at it in case there is a registrar that's trying to get away with doing nothing. Although those kinds of registrars don't hang out on this call. All right. We have Gabriel, and then we have Theo.

GABRIEL ANDREWS: Correct. Yeah. So, Paul, you summarized it great. And Martina, you did as well. And then Theo, I think you had some very constructive feedback, but I think your response was sort of parallel to the question, but not directly. I think maybe my explanation of the question wasn't making the most sense. So let's just, you know, as the example here. Picture a registrar has 10 accounts that are all the same bad actor and the same abusive scheme. And each of those 10 accounts has 10 domain names, right? Simple enough to imagine that.

And all the information is identical. All the same payment info, all the same information, et cetera. Theo, you mentioned that you could easily use email as your pivot point and identify all of them. And yes, great. That's what could be done. But what Martina was raising is, well, what if there is a registrar out there that wants to do malicious compliance, the absolute minimum check? Are we giving them the room with this policy to say, well, I used as my reliable internal data point, the one single account that was related to the one domain that I got to check the abuse report on, and then they find 10 domains instead of all 100.

Like, do we think that this policy stops that right now or enables that? And that's really the question I was flagging. So rather than saying at

least one reliable internal data point, you know, consider language that would address that sort of malicious compliance from the less responsible registrars that are not on this call, as Paul rightly called out. Thank you.

PAUL MCGRADY: Thanks, Gabriel. All right, Theo, go ahead.

THEO GEURTS: Yeah, that's a fair question. And it all boils down to how far do you go? I mean, when I do a reverse search on a database on an email address and I get 100 domain names from the same bad actor, then that is the result that I'm happy with.

Could I do more? That is a big question. Because let's play out two scenarios. I go with 100 domain names and decide that was enough and I'm fairly confident this guy doesn't have any more domain names. So report closed. It's completed.

Or I go to the scenario where I go look it up. I don't get any results. Now I go to check the phone number. That doesn't give me any results. How far down the line and within a wholesale registrar I don't go the line is very very short because then the fields that I have to my disposal will just run out. I mean it's currently tied to what the RAA says what we need to collect from data as a wholesale registrar; that is the data, I don't have any more, I don't have those IP addresses, I don't have the payment info, so I can't do a check on it, but the question really

becomes then, how far do I go again and again and again with each investigation?

And this is something that Volker already pointed out. If we are going to spend hours on these investigations, that will hinder all the other investigations. And then we will come into run into issues where we might overlook critical reports that are burning somewhere in some inbox, but we don't have the time to get into it.

And from a practical point of view, when we're talking about domain name association based on the information from the RAA, currently the email address does yield the most results. And if that is not the case any longer at some point, well, we will be reminded by all these other cyber investigators that are out there.

That's already happening. We already collaborate on various Slack channels like, oh, maybe we need to do this. Maybe we need to take that route. You know, these investigations keep on evolving. And yes, there will be a time that the email address will no longer be sufficient. In fact, I will predict within 10 years, this entire policy will fall apart because we need to move on. Thanks.

PAUL MCGRADY: Thanks, Theo. Most policies have a lifespan. That's for sure. All right. We have Ching, and then Martina. Ching, go ahead.

CHING CHIAO: Thank you, Paul. This is Ching from the BC. I think I would like to portray Theo's point. I mean, it's a well-said point, but I'd like to take it

in a kind of a different way to address it, because I think point number four we're talking about here, the one reliable internal data point, to utilize that, personally, I'm not thinking about how deep, how comprehensive it would be. I would use that as a kind of effective tool of how fast a registrar can help me or help others who are doing the typical, you know, anti-abuse work. You know, how fast, how efficiently this work can be done because of the registrar's ability to use that internal data point to discover an ADC.

The reason I'm saying this because of many vendors, including myself, my company, we've been doing bulk registration lists, DGA lists for a number of years based on different ways of, you know, algorithms, even sometimes guesswork. And then sometimes it gets delayed and delayed. It's not very timely. So the goal for this particular point on registrars relying on the internal data that they have exclusively is because of, from my point, is that to have the work done more effectively and timely. So thanks.

PAUL MCGRADY: Thanks, Ching. All right, Martina, go ahead.

MARTINA BARBERO: Thank you very much, Paul, and again, very quickly. I'm sorry for taking the floor again. I appreciate Theo's input on that, I think it's very helpful. But I think Theo is coming from the perspective of someone who is in good faith and wants to do a good investigation, and speaking about the depth and the length to which a good investigator would go. And I'm putting the helmet of someone who wants to avoid doing an

investigation in two clicks. Just say, you know, look, ICANN Compliance, I've done my best, this is the data point, and I found nothing.

So I think maybe the compromise or the way of looking at that is, instead of counting the number of data points, or prescribing which one needs to be checked, I think we need to give at least Compliance some ground to punish those that are playing the system.

And so, like, maybe instead of one reliable internal data point is one data point which is meant to bring to some connection or something like that explains that it's not just a tick the box exercise, but it's something that is supposed to be taken seriously, and that if Compliance sees that someone is playing the rogue game, then they can indeed try to sanction that, but yeah, that was just a thought that I think we can further discuss, maybe the GAC can come up with some text. I don't know.

PAUL MCGRADY: Thanks, Martina. All right. Theo, welcome back, and we'll have Farzi.

THEO GEURTS: Yeah, thanks. And to that point, I understand that there will be registrars who will try to wriggle out of these obligations. I do not think this will work. And why do I think that? When we change the contracts with ICANN and we enter language that we -- when there is a clear case of phishing, malware, botnets, and the other two that I won't mention, that created additional language for ICANN Compliance to take action against registrars and registries that were not compliant.

Now what we've seen from that change language in the contract is it did give ICANN Compliance a bigger bet to hit the registrar who were not compliant, but let me phrase it this way, but that bet is not big enough. But we've seen so far in the case of .top, it took forever till that registry became compliant. And of course, it was a bit of a unique case. It doesn't happen often that a cybersecurity firm just lodges a complaint with 200,000 domain names with boatloads of evidence. I understand that ICANN Compliance has to go through all the evidence. I understand that the registry had to go through all the evidence. But at the end of the day, ICANN Compliance got better and better at this.

However, with this policy that we are creating now, and with the language that we have here, it will give ICANN Compliance additional handles to go after these registrars who are not compliant. And with this language which we have here, they can do it faster, quicker, and more efficient. At least that is the prediction that I have, because I cannot imagine a world where all the cybersecurity, be it Infoblox, be it Netcraft, who's also here, I can't see a situation where Infoblox would report 20,000 domain names to me and I would simply ignore it. No.

First of all, I won't do that because we have a very good relationship with Infoblox and all the others. But I also notice in these Slack channels where we talk about registrars or where we talked about, and mostly those are not everybody's favorite registrars here; those are just the ones who are not doing anything. They get heavily criticized and they are already ramping up campaigns to report those registrars when this policy is finished. So I don't think there will be a likely scenario that somebody's gonna wriggle out of it. Sure, they will try it, but in the end

it will fail because this policy will strengthen the contract changes. Thanks.

PAUL MCGRADY: Thank you. All right, Farzi, go ahead.

FARZANEH BADIEI: I just wanted to say something briefly now that we are – like, we touched upon. 'Trust me' might not work, and we need to have better contractual language and commitment for all the registrars, and not for the good faith ones only. So we also keep saying that this is only limited to maliciously registered domain names and it's not related to compromised domain names. Like we need to also have languages there, guides, like strong languages that bind the registrar to only take action on a maliciously registered and have a strong language there.

And another thing that we keep hearing is, oh, this is all investigation and doesn't lead to action, like we are not prescribing action, so it doesn't lead to like takedown and stuff like that. That is also a 'trust me' argument that we need maybe a recommendation on how the investigation and associated domain check should not lead to mass and bulk domain name suspension. And we'd be happy to work on those languages as well if you're interested, if the group's interested. Thank you.

PAUL MCGRADY: Thanks, Farzi. And I think that goes to the compliance guidance, I think, rather than policy language itself. Because again, the more prescriptive

we are in the policy language, the less freedom ICANN Compliance has. But I'm open minded about that. If I'm viewing that wrongly, I think people should tell me. But I definitely think that those ideas are important to capture, Farzi. All right, we have Dennis. Go ahead.

DENNIS TAN: Thank you, Paul. This is Dennis, Registries. All this conversation about compliance, just wanted to add one data point. During the CP Summit that just wrapped up last week, ICANN Compliance made a presentation about their soon-to-be-launched proactive enforcement framework. And that is supplemented to the audits and other reviews, I mean, complaint-driven things that they do. I will post this live. And this is public information on the CP Summit website.

If you had questions about, you know, what is going the Compliance do, I think complementary to these mitigation actions that we are doing from a policy standpoint, ICANN Compliance is also stepping up into trying to find those registrars or registries that have the most concentrated reported abuse. So just wanted to add this data point in our conversation that, you know, this is not the only one.

This is one piece of the puzzle, right? That's what I'm saying from a policy standpoint, but there are going to be other PDPs for the foreseeable future that will complement ADC. But also, you know, on the other end, Compliance will step up a little bit more in trying to find those, as we say, willful blindness in terms of trying to mitigate. So, thank you.

PAUL MCGRADY: Thanks, Dennis. Okay, that seems to be the end of our queue. I'm going to do what I did last time and say, if there's anybody that has not spoken that would like to speak on this second topic before we move on, any stakeholder group, anybody that would like to say anything.

Okay, we have 16 minutes left, and my battery looks okay. So let's move on to some Deliberation over Charter Question 4. We'll throw 10, 12 minutes at this, and then we'll do some AOB and wrap up and get everybody out of here on time.

So Charter Question 4, what data access and privacy safeguards are necessary to protect both registrants and registrars during associated domain checks? All right. And so we have some preliminary language based upon the working group discussion on this one. An ADC must be conducted in compliance with applicable law, contractual requirements, and data privacy safeguards. Access to, use of, correlation of, retention of, and sharing of data in connection with an ADC must be targeted, evidence-based, and proportionate to the circumstances.

A registrar must not be required to collect or generate new data solely for the purpose of ADC or demonstrating compliance, and nothing in this policy should be understood to prohibit a registrar from using data, tools, or services that it may lawfully use in the normal course of its operations to investigate DNS Abuse.

A registrar conducting an ADC should limit access and correlation of data to what is reasonably appropriate for the investigation, taking into account the purpose of the check, the nature of the evidence, and the need to protect registrants and registrars. This policy should avoid

creating a prescriptive or obligation-heavy privacy framework and instead rely on high-level safeguards that can operate across different jurisdictions, business models, and technical environments.

The policy should not require incident-level sharing of precise information unless otherwise required by applicable law or contract, and any approach to retention of investigation results should be consistent with applicable law and existing contractual obligations rather than creating a separate retention regime under this policy.

All right. So this is a first stab at this language, right? It's not stable and nearly final. It's the opposite. And so we have looks like about seven or eight minutes to dig in on this one. And so let's take a queue. Let's look at the language. And if anybody has questions about the language, would like to see changes, think it misses the point of our prior conversation, this is the time. Any takers? Here we go. Theo, go ahead.

THEO GEURTS: Yeah. Well, I think the language already is sort of at a stage where it makes sense. I mean, it must be in compliance with applicable law and any other laws that are applicable in the sense I'm talking mostly about GDPR. So registrars, when they are not doing ADC at this moment, they will have to make preparations on access control. Who can do the investigation? How much information can they obtain. And these are all questions that registrars should figure out themselves. I don't think this is for this working group to sort of come up with new ways to comply with applicable laws. Thanks.

PAUL MCGRADY: Thanks, Theo. Ching, go ahead.

CHING CHIAO: Thank you, Paul. This is Ching from the BC. I think overall, I also agree that the proposed language looks good, except a small kind of idea here on the point of generating new data. So on this particular point, I'll probably also leave my notes on the Google Doc, is that I would hope that at least two data points can be generated. One is the ADC group number that the registrar is keeping track of. And then the ADC group size, how many ADC domains are in a particular group. So I'll actually leave my notes in the Google Doc, but just like to share thoughts first. Thank you.

PAUL MCGRADY: Thanks, Ching. And I get what you're saying, right? I mean, whenever we touch data, we create data, right? And so I think that it's not meant to be that you conduct your ADC check, but you're not required to keep notes, right? I don't think that's why. I think it was meant that they should not have to worry about creating new baseline data or whatever you want to call it.

We need to work on definitions a little bit around here, which was the purpose of the glossary. But after today's call, I'm realizing it more and more. And so I think that, yeah, your point's well taken, and staff will help me remember to give that some thought along with Nick and how

we make that distinction. Any other hands? I thought I saw Mary's go up briefly. Okay, Mary, go ahead.

MARY PENN: Can we clarify what we're protecting registrants and registrars from? So in that second full paragraph, we're saying that the ADC should limit access to, and then the need, we keep it in mind, the need to protect registrants and registrars. What's the aim here? What's the harm we are protecting against specifically?

PAUL MCGRADY: That's a good question. I have my own thoughts on that, but let's see if anybody would like to react to that. Anybody think that there's a possibility of a harm to a registrar or to a registrant? Theo, go ahead.

THEO GEURTS: Well, if you take the principle of GDPR and if you are processing data, you should have a legal basis. And this is what sort of is being suggested in the text here, the applicable law, again, being the keyword here. So you don't want to have a situation where registrars go off into a rabbit hole and start looking up all kinds of information about the registrants even if they are -- I mean I could do an ADC on a city. It will provide me all the registrants and the registrant info of everybody living in the city of Amsterdam.

That is something you shouldn't be doing in the first place, and I think this is reflected in the text that we have here, because you don't have

the legal basis for that, in my opinion, to do it. In terms of harm, yeah, that's a big question. I don't have the answer to that. Thanks.

PAUL MCGRADY: Thanks, Theo. I mean, I suppose one harm would be if you overdo it and you trip a wire under GDPR, that the DPAs will be unhappy with you. But I'll leave that to the advocates. Farzi, please go ahead.

FARZANEH BADIEI: So in general, data protection principles are to take action based on necessity and proportionality of the case at hand. And we have had scenarios where the registrar could actually be exposed to legal risk and political risk, if their investigation does not abide by data protection principles. I know that we don't have the time to go through the scenarios.

We are going to bring up the scenarios on the mailing list hopefully soon to talk about what sort of -- and we are not talking about like 100% harm is going to happen. What we're talking about here in the Data Protection Impact Assessment, as well as Human Rights Impact Assessment, is the likelihood of harm to the registrar and the registrant. And in our world at NCSG, we care about end users as well. So we make examples of that as well. So we will be clear on the mailing list.

PAUL MCGRADY: Thanks, Farzi. Look forward to that. All right, it seems like this is in pretty decent shape. But it is not -- good. Eberhard, go ahead.

EBERHARD LISSE: My view on this is, we should collect as much data as we need and keep as little as possible. Previously, there was a three-bucket list mentioned. For example, if the ADC yields nothing and otherwise we consider the thing innocent, then all the data that was collected must be destroyed or deleted. If a mitigation happens, the data should be kept for a reasonable time, something like a statute of litigation if you get sued or in any other way, and for everything else that is applicable local law.

PAUL MCGRADY: Thanks, Eberhard. And I would argue that even the statute of limitations and other document retention requirements are part of local law as well. So I think this is in pretty good shape. I think there may be some tweaks. This is going to go out on the list again in the collaboration document. I encourage anybody that would like to make those tweaks to do that within, you know, comments in the collaboration document.

I know Farzi's mentioned that they're working on a list of potential harms, which I suppose we need to in order to do some checks on these things. And so we'll look forward to seeing that. Again, this is not ready to be baked yet and be declared stable, but not final. And so we're going to have another reading of this. So anything that folks would like to see here, let's get that into the comment document.

We have five minutes left. We're going to go to Gabriel, and then we're going to move on to AOB. Unless there's somebody that hasn't spoken

yet that would like to speak, get yourself into Gabriel in the next 30 seconds. Gabe, go ahead.

GABRIEL ANDREWS: All right, okay. With regards to the last paragraph here, clarification question, where it reads, the policy should not require incident-level sharing of precise information. And in recollection of Rod's prior comments, which I would greatly appreciate with regards to the helpfulness and constructiveness of being able to share the abuse report that comes in.

For example, if a registrar receives an abuse report that has actionable evidence of how to locate additional accounts, I would love that the registrar be able to share that abuse report with any resellers that they're seeking to employ in additional checks. I wonder if this language as written would preclude that from happening. Over.

PAUL MCGRADY: Thanks, Gabriel. That's an interesting question. But we are going to make sure that we get out of here on time. Gabriel, if you can suggest some tweaks to deal with that concern in the collaboration document between now and our next call. Actually, ideally between now and Thursday morning, because that's when Nick and I meet with staff. I would appreciate that if you can.

All right. So we are not going to end up looking at the next thing. We're at AOB. I don't know of any AOB, but I am all about next steps. So keep looking at the strawpersons for Charter Questions 1 through 4. And

please go ahead and do the review of early inputs from the community that we have on Charter Questions 5 and 6.

And we are going to put out for a discussion and input something having to do with the earlier/initial impact assessments. So, that should be exciting. Stay tuned. And then we're going to do a staff update on the strawperson for charter questions 1 through 5, but it looks like 1 through 4 based upon the discussions to date so far. And so that's what we're going to be looking at. We have two minutes. Reg, 60 seconds of it all is luxuriously yours.

REG LEVY: Thank you so much. I would like to once again present an appeal to fix the chat. It is not currently useful because we cannot respond to things, so things get lost, and we can't react to things. So the chat is cluttered with a bunch of plus ones. I also can't italicize, I can't add page breaks. Anyway, please fix the chat.

PAUL MCGRADY: Thanks, Reg. Eberhard, 30 seconds.

EBERHARD LISSE: Of which I need 10 to figure out my computer. Now, remember, there is the Tech Day, and even while there may be a conflict, we are still interested to hear from the technical aspect of things. So if registrars or registries or even resellers have some stuff from a technical perspective that they want to present without creating a template for the bad guys, please get in touch with us.

PAUL MCGRADY: Thank you. All right, guys, we did it. I'm very proud of us. We did a good call. I always joke that we're going to set the land speed record in getting this PDP done, and I hope that's true. But there are times where you do have to slow down and dig in. And we slowed down and dug in. Anyways, we did it. Thank you, all. And we will talk to you next week. But please be active on the list and in the collaboration documents. I really appreciate the hard work today. Thank you.

[END OF TRANSCRIPTION]