

KRISTINE DORRAIN:

Thank you very much, Sue. Good morning, good afternoon, good evening everybody. This is the regularly scheduled meeting of the DAAR Working Group of the Registry Stakeholder Group. I hope everyone is well today. We are at a pretty exciting place in our working group. We've been doing this for a little over a year now. We've had a lot of meetings, a lot of discussions, we've produced some documents, some went somewhere, some didn't. We ran some hypotheses with OCTO, we checked those hypotheses. So, we have had a pretty full year.

And our last order of events at this point before we hopefully being to wrap up the group—we don't need to go on in perpetuity—is to do two things. One is to make a report of our findings to the full Registry Stakeholder Group, which we're planning to do on the next Registry Stakeholder Group call next week. And that's what we're going to talk about today. And it needs to be in a fairly final format, such as we all mostly agree on the concepts in it. Obviously as the working group has a chance to look at it, there's going to be feedback, so I don't know that we want to spend a ton of time wordsmithing unless we're worried about tone because I think ultimately when people get their fingers on it, it's not going to look exactly the same anyway.

Then the other piece, after we've showed our work to the Registry Stakeholder Group and gotten their approval, we are planning to offer the recommendations in the letter to OCTO. It would be on a call like this one where we would show them the recommendations, probably send them a written version of just the recommendations but not the findings or anything like that. And then ask them to kick off. And

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

there may be a little bit of additional work as far as looking at Samaneh's work or helping come up with some language. But really sitting around and brainstorming ways to make DAAR better. I think Jim and I have been talking that we think we've done a pretty good job for what we can do at the moment. And so, it feels good to point where we can show the registries what we've been doing and ask OCTO to run with the football a little bit.

JIM GALVIN:

With all sense of humility, of course, we did a great job.

KRISTINE DORRAIN:

Absolutely. We did an amazing job. We can all take a moment to pat ourselves on the back. So, with no further ado, let's dive in. I'm going to drop a link to the document in the chat just in case not everybody just got it on their email. And thank you very much, Sue, we have that up on the screen, and I'm going to just flip my other screen over here so that I can look at both screens while we're working. Okay.

So, most of you have seen this already. I'm not going to spend a lot of time going through the background, the summary of the work. This is work that you've seen before, you've had a chance to edit before, and this is work that you can absolutely wordsmith this if you want to. The summary really is just that bullet point of what I just talked about; the things that we've done and the things that we're proud of as a group.

Scrolling down to the top of the second page is where you'll see the first difference from our last call. So, originally what Jim had here was some

assumptions. And it worked out better that findings was really the best wording for what these were. This is the result of the work we had done and the thinking we had done and the factfinding that we had come up with during the course of our work. And so, it made more sense to present these as findings. And so, you'll notice that the findings are largely the same as the assumptions with some fleshing out. The sentence after the bolded word "finding" and the numeral is the actual finding and the words below it are really the background or the story about it. And so, as you read and as you have a chance to clean this up or think about it a little bit, that's really what it is.

So, for instance on Finding 1, the finding is that "Reputation lists represent varying views of reported abuse activity that is correlated with domain names." And then from there it elaborates on from that point.

So, you will see that this should not be super different from what you have seen before. I'm going to skip through the substance for just a minute as we look over the formal of the report and then we're going to come back. There are a couple of places where ... It looks like Rick's already got some comments in, so that's great. There's a couple of places where we had a few lingering issues within the working group that we need to resolve and some wording and some decisions that we had to make. And so, I think it's going to be ... I would like to spend most of our time there and have people take line edits offline. So, if you're skimming, feel free to make your line edits in suggesting mode, please make sure you're in suggesting mode. I'm going to flip that to myself as well just so that when I make comments, they show up in suggesting mode.

Finding 2, DAAR only shows the reported abuse activity. Finding 3, DAAR shows its calculated scores based on a point in time assessment. Finding 4, DAAR presents its reports as a new gTLD versus legacy TLD with no foundation for the distinction nor any useful conclusions. And Finding 5, there's an inconsistent understanding within the ICANN community about what DAAR is and what it shows.

And then we move on to the recommendations and these are essentially the same recommendations as before but just cleaned up and reorganized, taking into account the conversations we had last week and Kurt's email suggestion for how to reword things a little bit. You'll also notice that we don't differentiate between communication and the technical recommendations. And so, we sort of blur all of that together; it is one big set of recommendations here. We'll come back. What you'll notice here in Recommendation 4 is we reference Appendix C. That's the show your work. We talked last week about not wanting to lose track of all the "how" but not wanting to be super prescriptive to Samaneh and to OCTO. So, we move that into an appendix so the Registry Stakeholder Group can see the work and we'll have a capture of that work in the Registry Stakeholder Group's documentation and its history, but we won't be presenting that to OCTO unless it comes up in conversation.

And then we have a new section called "Potential Recommendations." These are the things that we think they may require more work, they require a bit more discussion as to whether or not we think they're a good idea. We definitely want the Registry Stakeholder Group to weigh in on the potential recommendations.

So, to be clear, we think recommendations one through seven—other than the wording and a couple of the features of those recommendations—are pretty well established. We think that this group has pretty good consensus on those. What we aren't sure about is whether what the consensus is around recommendations A through C, and so we want to spend a little time thinking about that today as well.

And then we just lobbed in a conclusion, which obviously could benefit from some wordsmithing. So, I just threw it in there so we'd have something to wordsmith.

The appendices are our first issues report that we prepared. That's Appendix A. Appendix B is the initial mark-up of the DAAR working group text that we sent to Samaneh that she said "Well, we can't really use it because we've changed DAAR so much, but it's a good starting point." And then Appendix C we said it's our homework from Recommendation #4.

Okay. That was a lot of talking. I need a sip of coffee. What I'd like to do now is turn to the substance. I'm going to pause and see if there's anything that gives anybody heartburn before we go to the comments in the margins. Otherwise, we're going to start with the comments in the margins as being the things that we want to jump to first. While I take a drink of water you think if there's something else you want to talk about first.

Okay. Very good. We will start with the comments in the margins. And the first one was just put in today by Rick under Finding #1. "Sometimes

a reputation list is primarily crowdsourced.” Rick is wondering if there are reputation lists that are not crowdsourced. I don’t know the answer to that. Jim, go ahead.

JIM GALVIN: Okay. I wasn’t planning on answering that question when I put my hand up.

KRISTINE DORRAIN: Sorry.

JIM GALVIN: But I can provide an answer there. I didn’t realize Rick had put the new comments up at the top, I should let you just run down the things. But anyway, to answer that particular question, what really happens is some of the reputation list providers run honey pots, okay? And they would ordinarily consider those more authoritative, if not absolutely authoritative, as compared to people who submit examples of spam to the—spam in particular the example—but they set up a honey pot to collect email messages in particular. So, that’s the way in which ... It’s not really crowdsourced, but they get an authoritative source for data. That’s my only example off the top of my head. I’d have to go back and look and see if something else jumps out.

KRISTINE DORRAIN: Okay. Thank you very much. And I’m going to keep your hand up then, Jim, because I think you were going to say something additional or

something different. I'm also going to take that point then that the correct wording isn't that sometimes a reputation list is primarily crowdsourced. I think then we would say "Reputation lists are primarily crowdsourced." I'm going to change it to that. You guys let me know if that's a problem because I feel like we want it to be pretty accurate and "primarily" doesn't mean "always." So, that does take into account the fact that there are other ways that reputation lists might get information. Let me know if that's a problem. Rick thinks that's an upgrade. Good. Jim, you'll go next, then Maxim.

JIM GALVIN:

You know what, let me withdraw my hand and let you go down through the comments. I was going to direct us in a different way but I cede the floor. Thanks.

KRISTINE DORRAIN:

Okay. Thanks, Jim. Maxim, go ahead.

MAXIM ALZOPA:

Actually, also some of those reputation lists might be ... They can have the source information for example from law enforcement. Some companies might use them, might collect malware and traffic from control—their bot control nets and try to get information out of the traffic and that software to understand where the control items are. Then they might add those domains. So, it could be ... We could add "and some other sources of information" to simplify that. Thanks.

KRISTINE DORRAIN: Thanks, Maxim. I would contend that all of the examples you just gave are still examples of crowdsourcing, just different crowds, right? Companies or law enforcement. So, unless you feel strongly about adding more language there, I think I would prefer to keep it as crowdsourcing because I think that keeps it at a higher level. It doesn't really matter who the crowds are. I see, you're thinking a closed network is not a crowd. Do you think it's important to say where that comes from? Do you think it's being more transparent if we say there's other ways?

MAXIM ALZOBA: I think yes, because it will show the whole picture without getting to unnecessary [details.] Thanks.

KRISTINE DORRAIN: All right. I'm going to make a note because I'm not quite sure how to word that right now. Who's running the closed networks?

MAXIM ALZOBA: I would say it depends. It's not official thing. It's a so-called network of people who are working on the same subject, yes?

KRISTINE DORRAIN: So, you're thinking like Shadowserver?

MAXIM ALZOPA: Yes. Primary. Shadowserver, they might have contacts with the crime agency with some—INTERPOL, yes? And for obvious reasons, they will never either disclose or tell anything about it because it puts in danger the operations, etc.

KRISTINE DORRAIN: Okay. That helps.

MAXIM ALZOPA: So, we should not name them, I think. Thanks.

KRISTINE DORRAIN: Okay. That's fine. I just wanted to get a sense of what you're talking about. Okay. I put it in bracketed, I'll fix it later. Okay, thank you.

Moving on to the next comment is Rick says "DAAR has its own weighted scale and is used in calculating scores. The scale may or may not be applicable to all users in all cases." Rick wants to know if the scale is consistent over time. Does ICANN provide notice when it changes? And Jim says there's no SLA, they have a published paper and it's probably reasonable to assume it won't change much but that's certainly no commitment.

Rick, any thoughts on that? I know there's kind of a question and an answer. Are you fine with that sentence as it is?

RICK WILHELM: Well, I think that maybe we should say that OCTO should be transparent about what the weighted scale is and then they should version it when it changes, right? So, maybe the weighted scale should be an appendix in the DAAR and it should say "This month's report was brought to you by the weighted scale version 1.2." And then when that changes, it is reflected in next month's report, that sort of a thing.

KRISTINE DORRAIN: Okay. Can we find—I don't know. Jim, you wrote the technical recommendations, is there a place to worm that recommendation into an additional recommendation that exists?

JIM GALVIN: Yeah. And actually, we'll figure that out. I'm thinking that it's either part of the messaging, that we want to fill out the messaging to cover this, or maybe even part of the community-wide webinars to explain it. We'll have to see. But, yeah, you went right to the same place I was thinking. Yeah, we need to expand a recommendation somehow. So, the comment is fine for now. We'll figure it out.

KRISTINE DORRAIN: All right. Great. I'm just leaving it in the brackets then so I'll come back and find it later. Beautiful. Carrying on. Rick says under Finding #2 "A great deal of abuse is the result of malefactors who take advantage of weak security on the service or website hosted at a particular domain name." I think it's the service or website host at a particular domain name. Rick, does that help?

RICK WILHELM: It was just a little ambiguous there when we say this service hosted a particular domain, just the way that the sentence was ... I know what we're trying to get at there. It's just sort of like a great deal of abuse is ... We're trying to get at the fact that abuse is taking place external to something that's controlled by the registrar or the registry and maybe this would be made more clear if we say it's weak security at the website or other service hosted ... By just inverting them and putting website first. That might make it more clear.

JIM GALVIN: Okay.

KRISTINE DORRAIN: Jim?

JIM GALVIN: Yep?

KRISTINE DORRAIN: All right. Was your hand up to respond to that or did you have an additional question or point?

JIM GALVIN: I think that's an old hand.

KRISTINE DORRAIN:

No worries. It's a very loosey-goosey format today as we try to fly through and wrap this up. Thanks for making that edit on the fly and, you guys, feel free if there are places where you think you can just fix the language without changing the meaning or at least without changing it controversially, feel free to do that.

Rick notes, still on Finding 2 at the bottom of page 2. "In addition, what is shown is the count of domain names that have correlated reports of abusive behavior. However, when mitigation is applied, a domain name is counted in one month and may not be present in the next month or may not have been removed from the reputation list." Jim added— Okay, never mind, I didn't see the comment he added. It looks like we handled that. Looks like we can resolve that. Rick, feel free to click "resolve" if you're happy with the way Jim did that.

RICK WILHELM:

Yeah, that looks fine. Thank you.

KRISTINE DORRAIN:

Great. Okay. Moving right ahead to Finding #4. It looks like everyone thinks Finding #3 is reasonable for the point in time that we're at.

Finding #4. "The DAAR presents its reports as a new gTLD versus legacy ..." Okay. "All TLDs should be treated equally from the point of view of having domain names potentially exhibiting abusive behavior." Okay. Seems to be a recommendation. Jim thinks it's interesting. I'm going to put a different viewpoint in. I don't know. I think it could go under either. The reason is because I view the findings as being

statements of fact that are largely incontrovertible and so I think that we should take it as a fact that TLDs should be treated equally from the point of view. Now, whether or not ICANN should treat them equally, that might be a recommendation. We might say ICANN's not and they should. But I think it's a fact that TLDs should be treated equally, but I'm waiting for someone to ... Rick, maybe I'll call on you. What do you think about my rebuttal?

RICK WILHELM:

So, really, I can agree with the statement, it's just more of where it's located, right? The finding is supposed to be a factual statement about ... I think that the way that we've been in this document structuring it is that the finding is a factual statement about the report, right? And the recommendations are a statement about how the report should be in a [2B] context, right? And so that's why when I looked at this would be both this statement and the one above where I'd highlighted a better approach, those to me—in here, in the document—we seem to slip toward recommendation language.

KRISTINE DORRAIN:

Okay. Yeah, all right.

RICK WILHELM:

So, both of those seem to slip toward recommendation ... Now for some reason, that one's highlighted on my screen but not on the shared screen.

KRISTINE DORRAIN: Oh, weird.

RICK WILHELM: I've got a comment in there where, for me, it says a better approach is highlighted and it says, "This seems to lean toward a recommendation." But that's not—did I not hit enter? Okay, maybe I didn't post that comment. There, you should see it.

KRISTINE DORRAIN: Yeah, there we go.

RICK WILHELM: I didn't hit post.

KRISTINE DORRAIN: Got it.

RICK WILHELM: There's another one that I didn't post there. Okay. So, sorry, I had typed those and then I forgot to hit "post."

KRISTINE DORRAIN: I always do ... I always forget that.

RICK WILHELM: I'm not very good at this whole Google thing. I didn't think it was going to catch on. What do I know?

So, both of those two sentences to me seem like they lean toward recommendations and where they might just need to move down into the lower section. But, I'm kind of shrugging my shoulders from a structural standpoint.

KRISTINE DORRAIN: Okay. We'll definitely take that on. I think I can definitely agree with you on the better approach one. On this one, I wonder almost ... I'm going to see if a friendly amendment changes your viewpoint on number four. What if we said, "The working group believes ..." So, that's the perspective that we are coming from as our finding rather than ...

RICK WILHELM: Yeah, okay. I'll buy that. And then later, I'll find that. Yep.

KRISTINE DORRAIN: Okay. All right. I'll just add that. And then what I think we'll do is as Jim and I do another read through, it may be that it also translates into a recommendation later. And our belief means that you should also be treating TLDs equally.

RICK WILHELM: Okay.

KRISTINE DORRAIN: All right. I think, have we then ... and I noticed you have a “they,” I think Jim fixed it so he’s fixing it as we [inaudible].

RICK WILHELM: Yep. That was good.

KRISTINE DORRAIN: Jim, resolve them as they go if you’d like. If you think they’re fixed, go ahead and resolve them. Okay. I think that’s it. Rick, you didn’t have any unsubmitted comments for five, right?

RICK WILHELM: No, I did not.

KRISTINE DORRAIN: Okay. Perfect. Jumping ahead to the recommendations. The comments that we have right now are the ones that we really left off with, and this is where I think I want to spend the bulk of our time. We do need to come to a conclusion in about 15 minutes here so we can go through the potential recommendations and go through next steps because we have a big meeting next week also. So, let’s dive into this issue about Recommendation 3 and I think it crosses over a bit to Recommendation 4 as well. So, it’s about this idea of a display of the top 10 or some number, right? We kind of discussed the idea of using this Pareto analysis to come to a number of abuse TLDs based on statistical analysis, not necessarily finding a hard list.

And not everybody agrees with the idea of naming TLDs. I put myself in that camp. I believe the other viewpoint is that eventually ICANN's going to get there, so we should put it out in a way that's more acceptable to us. So, I'd like to open up the floor to anybody who would like to recommend some language here about how to edit this and how to resolve this concern. Jim, you go ahead.

JIM GALVIN:

Thanks, Kristine. I think I want to say two things. One, just to add some context for people to keep in mind in this discussion of whether or not we should list any number of TLDs. I do think it's important to be aware that these lists exist. The reputation list providers maintain their own list of these things and their own overall stats and trends in this space. So, we're not hiding anything. The important message to take away there is we're not hiding anything by not allowing lists on the ICANN site, okay? This is an obfuscation opportunity for us, but it really doesn't buy us much and that's part of my personal motivation in suggesting that we should allow for the list to exist. I think that it's in our best interest to put down some guard rails or guidelines about how the lists are created and managed. I think that's in our better interest than trying to discourage listing because I just see that as inevitable given that they already exist. So that was one point.

Anyway, coming to the point about top 10—and I've been thinking a little bit about this, Kristine, and trying to respond to the Pareto idea that Kurt had laid on us. If we move in a direction of allowing for this recommendation, so if we move in a direction of allowing for a list of TLDs, then the issue under discussion here—the reason Pareto becomes

interesting—is, is 10 the right number, or not? Is there some other number which is more appropriate? And that's what Pareto gives to you. You go looking at the data and then there's just sort of this obvious cutoff and you know how many to list. And I'm thinking for the purposes of this recommendation if we get past the issue of choosing to allow for lists, I think all we have to do here is say "display the top TLDs" and just drop the number 10 and leave it at that, and we'll have discussions when we iterate on this as we get to the conclusion and we talk about the next steps. The goal here is to be iterating with OCTO on exactly what they do. That will be our opportunity to suggest the Pareto analysis, which I'm sure Samaneh well understands. And assuming we can get some agreement on all of that, then the right thing will happen. There will be an appropriate number of the top ones and the long tail business will just be left off. And whether it's three or 10, the right thing will happen. So, thanks.

KRISTINE DORRAIN:

Thanks, Jim. And I think I'll offer the alternative viewpoint and I think Jim and I are both pretty well-established in our positions here. But the view that I have is the lists do exist, but that doesn't mean we have to ratify them or we have to volunteer to name and shame one another. The Registry Stakeholder Group has always stood for not necessarily closing ranks but even if we disagree with one another we just move along and do the best we can around it. And so, I don't see the Registry Stakeholder Group being the first ones to go to ICANN and say "We should list the worst among us on your website because we think that's the right thing to do." It's not that it won't get there and it's not that other groups aren't going to recommend it and it's not that ICANN

won't do it. It's that I don't think it's the position of the Registry Stakeholder Group, whose mission is to advocate for registries, to essentially throw one another under the bus, even if some of us deserve to be thrown under the bus. And we've been really clear about that in our public comments and in the way we just approach controversial issues.

And so, I don't believe the Registry Stakeholder Group is going to support a list and so my concern with putting this out there is one, it's not a consensus view of this group and two, I don't see the Registry Stakeholder Group overriding that and suddenly jumping on board and making it a consensus viewpoint of the stakeholder group. And so, I'm worried it's going to take the entire conversation away from the amazing work we have done and the fantastic recommendations we have here and it's going to put all of the focus just on that.

Remember, not everyone on the list is as evolved as we are with respect to DAAR. There's a lot of people in the stakeholder group that still hate it. I'm one of them. But there's a lot of people that still hate it. And so, I think the idea is I would hate to have all of our work get overshadowed by this one point.

But I would love to hear from other people. Jim and I have each articulated our positions multiple times. We absolutely respect one another's positions, we absolutely see each other's sides, I mean there's no danger of that. But those of you who hadn't raised your hand lately, I would love to get a viewpoint from some of you as far as what you think about this point. Sam, go ahead.

SAM DEMETRIOU:

Thanks, Kristine. I think what I'm struggling with a little bit, and I think you alluded to this in what you just said, there seems to be an inherent contradiction in my mind between raising concerns about the feeds and their accuracy and the way data is presented, but then also saying "Oh, but go ahead and name the TLDs that, according to your metrics, are 'the worst,'" right? I think when you look at this document and these set of recommendations as a whole, it seems like those two points are a little bit at odds. So, I don't know if maybe— understanding that there is a division of viewpoints, right? And I don't know where other members of the stakeholder group would fall on this—maybe this is something we think about in phases. As in, make the improvements that we suggest up front and then later reconsider the question of whether TLDs should be named if appropriate improvements are made. Not committing to it, but understanding that this could be something that's addressed at a later point in time.

Just trying to think of this in a way that knits together the different viewpoints we have here. But I definitely agree with your point that if we drop this on members right away, it is probably going to distract from the whole conversation.

KRISTINE DORRAIN:

Thanks for your viewpoint, Sam. It looks like Beth has a comment in the chat. "I think that makes sense," I think she's referring to what Sam just said. "If more confidence is gained following the improvements to reporting, then we should reconsider." And it looks like Jim is moving in

that same direction also. He's got a comment that says "Remove to propose recommendations." And I think that seems right. I think that we can absolutely frame that as a discussion that we've really thought hard about, because we have. And I think any one of us here are proud of our efforts and we're here because we're proud of our efforts. And so, I think that wanting to showcase our pride in those efforts is certainly admirable. Does everybody then agree that we should move Recommendation 3 to the proposed recommendations? Rick, go ahead.

RICK WILHELM:

Yeah. I certainly agree with the point that Sam is making about the fact that having us bringing up the issues with the feeds and then also simultaneously recommending we name names would be inherently contradictory. And so, I don't know if having them in a ... Okay, sorry. You had call them proposed recommendations, but then I scrolled down and it's called potential recommendations and maybe calling it potential recommendations is a better spot. The point you had made toward the beginning of your position statement, Kristine, I thought was very valid and very, dare I say, moving about the notion of the stakeholder group being engaged in self-immolation and throwing a certain of its number out, hauling them before the court of public opinion, seems very odd to me for a stakeholder group to be putting that forward as a recommendation.

So, I think we ought to still consider that point of view very carefully because no matter how the feeds get going and get improved, the notion that we would be reporting that stuff externally seems surprising to me because I thought you stated it very well. I think that the

improved feeds would certainly give the registries and the registrars more ability to take action if the feeds do get improved, but as we've stated before, there's a lot of limitations in those feeds. And those folks that engage in those feeds on a deep basis right now, in a tactical manner, know that names are slow to come off of those things and make it really hard to take action based on those things. Thank you.

KRISTINE DORRAIN: Thanks, Rick. Beth, I see your hand. Go ahead.

BETH BACON: Thanks, guys. And I'm going to apologize in advance, I have a couple of questions. I have missed the last call or two and I apologize. So, maybe this question is out of date. With regards to potential recommendations, I'm assuming that means potential future recommendations. So, my thought is that if we aren't going to include it in the recommendations, we could keep it out of the document that we give ICANN with an understand that we'll come back and we'll revisit this as ICANN either implements or doesn't implement our recommendations and we'll review this again in the future. And this could just be an issue that they propose again or that we bring up and say "Sure, we think that this has matured to a level of confidence where we could display some things." Again, Kristine, I do think that your evaluation of the Registry Stakeholder position is accurate and it's going to be a perhaps nay to a hard no.

If we do want to keep this in the document as a potential nod to it as a kindness, a little olive branch to ICANN, I would suggest that we would

reword this and make it less specific and move it to a potential future recommendation as opposed to just a potential recommendation because my understanding is the potential recommendations are not really being made, they're just items for future consideration.

So, if we did move it to a potential future recommendation, I would say let's reword it a little bit, make it less prescriptive to the top 10 TLDs. We could just say for each type of abuse activity in the future we will consider some level of publication of statistics or percentage of domain names, something a little more generic. So, I think that's it.

KRISTINE DORRAIN:

Thanks, Beth. I really appreciate it. Yeah, and Jim has answered your question in the chat, but I agree completely about the potential recommendation, that's exactly what it's supposed to be. So, it's meant to be a bit of a parking lot, right? So, these are issues that in Amazon we call them considered and rejected, or considered and tabled. Basically it's the stuff we thought about and worked on but it doesn't rise to the level of a current recommendation. So, as Jim mentions in the chat, we're going to present it to the Registry Stakeholder Group. We want it to be as part of the complete documentation so those that come behind us five years from now or whatever can see what work we've done and thought about.

It may be that on the call next week somebody ... There's a mass movement behind—I'm going to just randomly pick one—Recommendation BA. Everybody says "that's the one we ... Why didn't you include that? We like it." I'm just randomly picking it, so we don't

have to read it. But we could decide in the moment that we're going to pull that up into the recommendations, absolutely. And so, when we talk about them, they really are for cogitation and people saying "Yeah, we should do" it or not. But it's also meant to be that parking lot for the future. So, I think your suggestion is absolutely right.

I wanted to offer in a friendly amendment as well as we're talking about the wording here. And I like the words that I heard people say. I think both you and Sam talked about it a little bit, but how we think about these lists maturing and our trust in them growing. And we could say in this recommendation after we move it to the potential future recommendation section ...

Once the Registry Stakeholder Group feels confident in the changes in DAAR such that we have a higher degree of confidence, the lists themselves have matured, the ability to demonstrate the persistence, which we're going to talk about in Recommendation 4, is demonstrated, once we actually can quantify that these other recommendations have been accepted and implemented and are making changes, then at that point we think that we should look back into this—what's currently listed as Recommendation 3—so that we can revisit it and decide, is now the right time to make this next leap? So, I think that's right.

And Sam, to answer your question, in my hand the potential recs are up for discussion with Registry Stakeholder Group members, but if people don't agree or support, they would not go in the final doc to ICANN. That is correct. Whatever goes in with the potential recommendation section is just for us to discuss. If the group feels incredibly strongly and we want to pull one up into the recommendations, we can. But they're

the ones that we don't think people will be ready to pull up into, which is why we've separated them out. We don't want them to bog down the rest of the meaningful discussion around the recommendations that we can all whole-heartedly agree on, which I think is one through seven, minus three right now.

Beth was going to ask that same question. Yay, okay we got two birds with one stone. Jim is adding some friendly amending language—or somebody is, it looks like Jim—into Recommendation 3 and it looks like we're going to move it to the potential recommendations section. I think that's very helpful. Thanks everybody.

All right. Now let's move to the big, exciting recommendations, the ones that I think are the most new and interesting when we talk about this persistence of reported abuse activity. That's the new big idea, right? This is the one where we actually get to show our mitigation and show our work and demonstrate that we're good actors. And I think it's as a result of Recommendation 4 that we will probably get to Recommendation 3 at some point.

So, the unfortunateness is we moved all the beauty of Recommendation 4 into Appendix C, which is all of the math and the homework. But the next comment there—I'm going to scroll back up so I can see what the comment says. Rick says "Let's discuss this recommendation. It's been discussed how names are sticky to get on that list." So, Rick, I'm not following your comment. I'm going to offer you the opportunity to explain. Go ahead, Rick.

RICK WILHELM: Sure. I'll comment here and maybe Sean might comment also—just to prep him for that. I think that this is going to be interesting where if OCTO does agree to this and actually does do some measure of the persistence of report activity, it's going to show some interesting behavior where it shows that names get on the list and stay on the list despite mitigations within the registration ecosystem or within the service provider ecosystem. And it shows that the lists are long [lived], if you will, and there's burn-in on the list. And I think that it's going to show that and that it's hard-timed ... It's sort of like for some people how they have a hard time getting a ding off of their credit report, metaphorically. That sort of a thing. I think that's what it's going to show and I'd welcome comments from Sean. I see Jim's hand up also, but I know that Sean has commented in this regard and other times also, but maybe he can comment in that regard. Thank you.

KRISTINE DORRAIN: Thanks, Rick. Sean, go ahead, and then Jim.

SEAN BASERI: Okay, great. I think what Rick mentioned will likely be the case. And I'm wondering if it might be of value to consider this after we have a little bit more information, because one of the things that happens is these data providers will collect their information, make their assessments, produce their data feeds, and those are used in a lot of different ways. And as Jim has mentioned many times, it's used in enterprise-level internet filtering devices or firewall systems or email systems.

But I'm wondering if this would be more valuable as a recommendation once we have a little bit more information about how some of these data feeds are listing domain names and then removing, their criteria—at least the removing part, because as Rick mentioned, there may be situations where some have time periods where they'll just unlist them. Then the question becomes, if this is published, is it really an accurate representation of the work done by registries, for example, if no matter what's done, a name is removed only after a period of time?

KRISTINE DORRAIN:

Thanks a lot, Sean. I think Jim's got a couple of answers to that and he probably has his own comments as well. So, I'll let Jim go and then I'll fill in.

JIM GALVIN:

Thanks, Kristine. I agree with Rick and Sean and I completely understand the point that they're getting at. It is true, well maybe more of an observation, that reputable reputation list providers, if you will, all have documented procedures for how they manage their list and mechanisms do exist for individual domain owners to get on or off a list. Sometimes those lists are managed in a relatively automated way so that if whatever behavior was bad was happening stops, then you'll drop off the lists. And all of that is true. So, again, reputable reputation providers, they have feedback mechanisms and they do work with those.

So, from my point of view I'm after two things in why this recommendation is here and how I've been trying to characterize

something which I think would be useful for us. And I apologize for maybe talking a little bit long here about this, but there are two points in my mind.

One point is exactly the point that Rick and Sean are making, which is that if we do this, it actually should highlight, in fact, the quality of some of the reputation list providers, or more precisely it will actually highlight the quality of the feedback mechanism that exists in the wild, if you will. And that might be something that we can take advantage of at some point in the future. But, in any case, that might be something that we ought to add in our discussion here or maybe as we're iterating with OCTO later on in our next steps here, we need to get them to revise their messaging—to call out this fact. That's what all of that means. That this persistence in reported abuse activity has limitations. Because you're right, we don't currently talk about the limitations of some of these metrics. And that might be something for us to consider.

So, from my point of view I guess, I actually like that we're going to highlight that particular quality issue. Others may have a different opinion and that's lovely and I mean that quite sincerely. If we have a different consensus here in the group, I'm not going to die on this sword or anything, but let's just keep that in mind.

The other thing that I wanted to say is we've had a few discussions about how to show persistence and how to show that we get mitigation credit. This is simply one suggestion for how registries in particular would be able to show some amount of mitigation credit. It's not a perfect proposal by far, and we've already highlighted one particular

flaw in it, but it has one particular feature that I like. And I want to call it out here just for the purposes of discussion.

The important feature here is there's no work on the part of registries. In other words, just showing this is something that ICANN entirely has to do and then they put it up and they do that. And it's just taking advantage of what's out there in the system, [and it's present and it exists.] We have talked about a couple of other ideas for how we might show that mitigation is taking effect. I won't go back over them now, I just observed that all of those other mechanisms—there's only been two, or in fact I can only think of a couple offhand, but they all require registries to do something. They all require an obligation on us to take some kind of action that can then be incorporated into DAAR. One of the reasons why I like this mechanism is not just that it highlights a flaw overall, it contributes to the overall flaw in the DAAR system, which I see as a feature, but it also doesn't obligate us to anything and I was regarding that as a feature, too.

Now, if we want to view this differently and look for better quality, then we have to be prepared to look for us contributing to the work and [doing that happen] because I really do believe that one of the requirements that we have in this group is we have got to find a way to give credit for mitigation. And that's the way I'm phrasing it, as a conceptual concept. Somehow, that has to be incorporated in DAAR and this is just the leading proposal at the moment, although certainly, other options are available if we want to consider them. So, I guess that's my comment. Thanks.

KRISTINE DORRAIN:

Thank you, Jim. A couple of comments before I continue in the queue. One is that I was just going to echo what you were saying about this is all on ICANN, the feeds don't have to get better or smarter themselves for this persistence program to work. ICANN is supposed to be doing this work and, if anything, one of the features I like about Recommendation, now 3, formerly 4, is that it will actually demonstrate the failure of these feeds because ICANN will show—this is I think Sean's point—that they're listing domains that should no longer be part of their lists. So, it's one of the things I'm hoping here.

I'm going to close the queue after Beth. We want to try and wrap up. If everybody would take no more than a minute or so, that'll give us a few minutes just to circle back on what's going to happen next week. Thank you. Go ahead, Rick, you're next. I think, Sean, that's an old hand.

RICK WILHELM:

I'll be very quick. I just want to point out that while I was talking about these lists having burn-in and being sticky on, Sean brought up a very important point that I think it's important to realize that in the commercial sense in which these feeds are used typically and people pay for and subscribe to these feeds, this burn-in is actually a feature, not a bug, because people want these names to be sticky on so they can protect their enterprises. Because one of the biggest indicator that a domain is going to be bad in the future is if the domain was bad in the past. And so, the reason that they are sticky on is because the domain is likely to be bad in the future. So, it's actually not something that the reputation lists would actually define as a bug or a failure in their product. We, in the registry ecosystem who are trying to minimize those

names being attributed to our registries or registrars, would see it as that, but in the RBL business, they would actually say “No, that shows the effectiveness of our product.” Thank you. And thanks to Sean for bringing that up.

KRISTINE DORRAIN: Thanks. I hope I captured at least the gist of that in the comments. We can hopefully work that in. And Beth, you’ve got the final word.

BETH BACON: Thanks. I just wanted to say I appreciate all the context that Rick and Sean and Jim have given. I think the important thing about Recommendation, now, 3, is that we explain pretty thoroughly how it’s different than if it’s displaying the persistent reported abuse activity, ICANN has to be pretty clear about what the parameters of that is if we are concerned about the stickiness influencing that. Because if we don’t and we’re not clear that it’s not influenced by the stickiness or it’s not going to negatively impact us simply because of the stickiness—I don’t ever want to say that word again—then I don’t see the difference between that and what recommendation formerly known as three is. It would just be another mechanism to say these folks are doing badly, but it would be inaccurate. So, I just want to make sure that we are very clear with ICANN with what that means.

UNIDENTIFIED FEMALE: Thanks.

BETH BACON: And it sounds like we are.

KRISTINE DORRAIN: All right, thank you so much Beth. We'll close the queue for now. Here's where we're at. Jim and I will take these comments on board, we'll go one more run through hopefully today. But you all have this link. Please turn on suggesting mode, make your suggestions and your comments or ask your questions so we can respond to them in the doc. This is our chance to get that doc cleaned up and prettied for the Registry Stakeholder Group. We'll have another call to talk about what the doc to OCTO looks like and I think it's just literally going to be a copy and paste of the recommendations and that's it. But that's going to be up to the group at that time.

Next week we will do a quick run-through of any big, sticky points. We will not do a line-by-line review like we did today. Literally if anyone has a big burning hairy question in the doc, we'll bring it up. Other than that, we're going to close out the doc first thing next call. Then we're going to go through the PowerPoint, I'm going to put together a PowerPoint to get ready for the Registry Stakeholder Group call where we'll present this information. So, we will send the doc out right after our call on Tuesday to the Registry Stakeholder Group so they've got 24 hours to look at it. We will also have a PowerPoint up so people don't have to be reading the whole doc on the call. And then we'll be able to answer questions about the doc itself.

And so, next week what we really want to do is anticipate those questions, think about the messaging, and obviously we are deep into

this, what do we need to provide for background. And once I hear from the ExCom with Donna and Sam and Beth what timing we have—I don't know how much time we're going to have also, so we're going to have to work within the parameters of the time allotted.

The final thing I wanted to leave everybody with, and I think probably most of you know this, is that I am leaving my role at Amazon Registry. And so, I will be stepping down as co-chair after next week. Next week I'll be helping put the PowerPoint together and give the presentation to the registries. My last day with Amazon Registry is actually Friday, but I'm sticking around to help you all out one more week. And then it'll be up to you as a group what you guys decide to do about another co-chair. But I'm hoping that we're leaving with most of this work done and that may or may not be necessary.

So, I wanted just to say that this is a sad moment for me. If I'm not connected with you on social media in some way, please feel free to connect with me. I want to stay friends with everybody. I'm going to be staying in the abuse community. My new role at Amazon is senior corporate council for content policy, so really helping Amazon come up with internal policies for making sure that bad guys don't do bad things with the content on our services. So, I'm going to stay in the abuse world but just on not the technical side.

Anyway, I wanted to say thank you to everybody. This has been one of the most fun things I've done in the Registry Stakeholder Group and so thanks for giving me the opportunity to work on this project with you all. And I'm not going anywhere, I still have the same email address if

you have any questions. But I'll be around next week, but I just wanted you all to know that's where I'm at right now.

With that, I think thank you very much to everybody. And please turn on suggesting mode and edit the doc. We'll see you all next Tuesday. Have a great day. Sue, we can end the call.

[END OF TRANSCRIPTION]