

---

KIMBERLY CARLSON: Thank you and welcome to today's NCAP Discussion Group Teleconference on April 29 at 19:00 UTC.

In the interest of time, there will be no roll call. Attendance will be taken based on those on Zoom. Kathy and I will update the wiki with the names of the participants as quickly as possible. We do have one apology from Russ Mundy. All calls are recorded and transcribed, and the recording transcripts will be posted on the public wiki. As a reminder, to avoid background noise while others are speaking, please mute your phones and microphones. And with that, I'll turn the call over to Matt, I think.

MATTHEW THOMAS: Sure. Thank you, Kim, for that. Good afternoon, everyone. I hope you're all doing well and staying healthy still. We can go ahead and get right in to the agenda.

Does anyone have an update to their SOI at this time? I see no hands, so I'll take that as a no. Again, we have any new members this week, but again I'm encouraging you all to reach out to folks if they have interest to come join the discussion group and help us through this. I think the big thing that happened was that since our last meeting, Matt Larson has sent out the Phase 1 final report. Matt Larson, would you like to say a few words about that now?

---

*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.*

---

MATT LARSON: Well, I could. I'm not sure what I would have to add, to be honest, though, to the exchanges we've already had in e-mail. I could be very brief. I'll say that I've requested anyone's feedback and questions by Friday because OCTO has committed to the Board to have the report and our Board document to go with it to the Board by the end of June. We have a public comment required in the project plan between now and then, and that is a minimum of 40 days that the public comment itself has to be opened plus time to prepare the report. So by the time you add all that up, we're going to be down to the wire unless we allow Karen some time to process this group's feedback, make edits, and then get it out to a public comment. That's the reason for the one-week review period.

MATTHEW THOMAS: Thanks for that, reminder, Matt.

MATT LARSON: I guess I should ask. If anybody feels that they need longer, I could consider it but I have said to the Board Technical Committee multiple times, I keep promising June 30. I need to do what we need to do to make sure that we make that day or it's going to make me and the rest of OCTO look not good, but I'm willing to listen to request for slightly more time if that is absolutely necessary.

MATTHEW THOMAS: Thanks, Matt. I hope most people have had an opportunity to take a look at the document, read it through, see the new changes, specifically

---

Section I think 5 and 6 were most of where the new work was. Do you want to have any kind of discussion about the document right now? Yes, Jim?

JIM GALVIN: Yes. Thanks. You can hear me, right?

MATTHEW THOMAS: Yes.

JIM GALVIN: Okay, good. Jim Galvin from Afilius for the record. Yes, speaking a bit as co-chair for the moment, chair. I do want to have a discussion here about this document. We've obviously had a bit of discussion already on the mailing list. There is one question which really does need to be answered here, but maybe let me try to put a little bit of this in context. Some of this has been said on the list and so we'll just try to package all this up and hopefully people will track all of this.

There is an important distinction to be made between whether or not there is additional funding for Studies 2 and 3 with respect to the Board project and the original project that was proposed by SSAC to the Board more than two years ago at this point. That is sort of one half of things. And the document that Karen has produced under a project that was managed by OCTO is the first part of all of that and they produced there their work product which has a particular recommendation about whether or not to fund future activities based on the work that was requested then at that time.

---

It is important to recognize that as a distinct thing from whether or not the NCAP project, NCAP Discussion Group continues. SSAC has been tasked by the Board to answer a set of questions and to provide a response to that. The process by which SSAC is going to do that is we created this NCAP Discussion Group in order to meet one of the Board requirements which was an open and inclusive process. And so we created this project and this discussion group in order for that work to continue as an ICANN open and inclusive working group. That's what we're doing here. So whether or not the Board funds future actual studies as originally defined by the project proposal that had been put forward by SSAC is an important consideration but it will not stop this working group's work. This working group and discussion group will continue because we will need to do the best that we can to answer the Board's questions and provide a response to the Board. SSAC has that obligation regardless, and so that work will go forward.

Now, a question that comes before this discussion group at this time is that given that the report that the work product that was produced by Karen and she properly made her own assessment, which is fine, we have to recognize that that was her assessment of the details that she saw, we did produce for ourselves this discussion group a technical gap analysis for ourselves about what we thought would be useful or is important to recognize as far as going forward with Study 2 and Study 3, the analysis part, essentially, the overall NCAP project that originally was proposed.

So it is reasonable for this discussion group to ask itself, do we want to respond to the recommendation that OCTO is going to put forward that as produced by Karen as produced for the work product one? Does this

---

working group want to respond to that with its own advice and recommendations regarding Studies 2 and 3? That is a question that we get asked ourselves. Do we think that we really do want some funding to proceed with the technical gap analysis that we had produced to ourselves a couple of weeks ago? Do we think that that's a relevant thing and a basis by which we should ask for funding? We get to talk about that and make that decision here and then produce such a recommendation, and SSAC has its own mechanism for providing that advice to the Board if that's something we want to do. We might also choose to just want to respond in the public comment period. The final report will have a bit of a public comment period, we could also respond in to that with respect to the document. That is certainly something we could do, so there's two paths for us to take and opinion if we choose to do that.

There has been some discussion on list about some questions and there's been some answers to some of those questions. I think I'm going to pause there at this point. I just wanted to focus on the logistics and process, ask first if anyone has any questions about that, and then open the floor for people want to raise an issue. I would like to have opinions and in fact substance to the question of, should we have something to say? Should this discussion group have something to say about the conclusion that was created? I guess, if you don't mind, Matt, I'll continue and sort of manage the discussion here at this point and turn it back over to you when we come off of this. Jeff, go ahead, please.

---

JEFF SCHMIDT:

Hey, Jim. I put this in the chat but it sounded like what you were saying was these were competing papers, if you will. We had to support the technical gap analysis document or support the, for lack of better term, the Karen document. I don't view it as an "either/or" like that but that does raise the question, what is ... I'm not sure what the practical outcome is. Both documents are correct in my opinion but they lead to different outcomes and they answer different questions. So, sorry, that's not a very helpful question, more of a comment. I'm not sure exactly what to do with.

JIM GALVIN:

Thank you for asking for some clarity. Let me try to frame this a little bit differently. Like you, I really don't see these things as competitive. I mean, they had a task and a project that was created and they produced the work product based on that request. In the same way, we have our own work product that we have to create and we have to get to. So the question here is whether or not our technical gap analysis is competitive with the report produced by Karen and her recommendation. If anything, I would think that ... it actually is a broader scope. I actually think that our technical gap analysis said some additional things that were not covered by what Karen said. She wasn't asked to produce comments on everything that we produce to comment on. So I would view it as a constructive addition to the advice that she's providing forward to the Board.

The question that we need to ask ourselves is do we think it's important that we get the additional funding to get an additional contractor to do some analysis work that we think we need? Is that important to us in

---

order to move forward? That opportunity still exists. Our technical gap analysis brief that we created for ourselves provides some substance to support that request of the Board and we have a means to provide that to them. Or do we just want to let OCTO submit its report, wait and see what the Board says, keeping in mind that we're going to go forward regardless. Either we're going to get some additional funding to do some additional detailed data analysis or if we're not, if we want it, we should ask for it. If we're okay with whatever that works out, fine. Then we'll just continue on with what we're doing. That's the question before us.

I hope that helps, what you were asking for, Jeff. If you don't raise your hand up to want to speak again, I'll assume that you're okay with that. Any questions from anyone else, or comments? I hope that that's clear.

The folks that were making comments on the mailing list in particular, do you want to speak up? I guess I'm looking primarily at Danny and Jeff Neuman, in this case, were the principle ones who commented on the list. Do you want to bring anything out here and offer an opinion in response to the question of having our own advice to the Board with respect to moving forward? I'm not seeing any hands and no one wants to jump in. Although interestingly, Danny moved to the top of the queue here, which suggests you raised a hand, but I don't ... I'm not hearing you, Danny, if you're talking and I don't see a hand.

DANNY MCPHERSON: Hey, can you hear me?

---

JIM GALVIN: Yes, we can. Go ahead, please.

DANNY MCPHERSON: I don't know which microphone I'm using. I'm sorry about that. I was going to raise my hand as well.

I think the operative phrase in one of Karen's comments was the study is currently designed. So maybe we just take a look at that. I thought all along that I wasn't sure that we could torture the data and get more answers but I thought that there were other strategies where we could more intelligently address things than the brute force method that is controlled interruption. And I think that Jeff Schmidt's paper he submitted to the mailing list a few minutes ago was a testament to that saying people should probably change their configurations that are resulting in collisions. We know some examples that have been seen now that ICANN reported on from people working from home, things much broader than corp.com. I'm not sure if Jeff wants to comment about the thoughts of controlled interruption across the board for collision mitigation. But in response to his paper or in response to the corp.com sale, I would be like, why don't we just apply controlled interruption? Except that, as Anne pointed out in her e-mail response to Jeff Neuman, there's a whole class of attacks that you never see at the brute-force level at the registry which are man-in-the-middle attacks that could occur. We've seen blueprints for these on mailing list and other things. Somebody wants is a man-in-the-middle attack and somebody gets hacked as a result of such an attack then nobody is reporting that to ICANN. So I don't know what it is that anyone thinks would be reported to ICANN when even the victim doesn't know they're



---

a victim. So I think there's a whole class of things there that need to be further explored. I do personally intend to provide some written comments in response to the conclusions of the report.

The other thing I would say, by the way, is for the record. My comments on the mailing list to Karen, in response to this that Matt Larson replied to, were about Karen's research and findings. What I intended to convey with my quotes was that this was the working group's research and I think the report itself is quite comprehensive. I just don't understand how it reaches the conclusion that it does. So that's what I'm going to look for.

With that, I'll answer your actual question, Jim, and say that if OCTO isn't even commenting on Karen's comments yet, then I don't know how NCAP could comment without OCTO's input, and I think that SSAC is going to provide a comment clean of both those things. So I am a little surprised, there are three or four additional stages of comments and documents that are going to be developed here, but at the same time I think even where we're at, there's not really any alternative. So I'll stop there.

JIM GALVIN:

Thank you, Danny. I want to go back to one thing that you said in the very beginning where you called out the phrase "as designed". I agree with everything else that you said and I don't want to highlight this in particular for the rest of the discussion group here in case they missed it. It is true that the original NCAP proposal was designed in a certain way and set up certain questions and certain things that it wanted to

---

do. I do think it's useful to recognize that that was two years ago. In Internet time, two years is kind of a long time. No matter how you look at it, no matter how you cut it, it really is.

In that respect, with Karen's observation and the document she had produced that the proposals for Study 2 and Study 3 as designed, it's entirely possible that it's okay to view them as not really the right way to approach this and the right thing to do. I think there are different sets of questions. In fact, I would make the observation that our technical gap brief that we had created for ourselves, there's obviously some new information and we highlighted the desire to do another set of analysis. Even if it's a repeat of what was done in 2012, we had suggested that it's important to do that because the infrastructure is different now. At a minimum, the technologies that are deployed that exist in the world are different now than they were then. We just think that it's useful to look at whether or not that changed the traffic in any way, that that changed traffic patterns. It's entirely possible that nothing will change and it will all look good. But I count that as ... that's a success to know that those particular technologies did not have any impact on the data that would ordinarily be available. I just think that that's an important question to ask. It's also useful to recognize, as you said, with as designed, there probably are new questions and that's part of what we have to get to, we the discussion group, in our beginning of our Study 2 analysis that we have to get to what kinds of new questions do we think they want to ask to help us. Now we've got quite a few forming here. So let's get some other voices. Jeff Neuman, go ahead, please.

---

JEFF NEUMAN:

Thanks, Jim. I caught some of Danny's comments. Sorry, I actually got disconnected from the Internet for about 10 minutes. Well, I'm back. So, Danny, if you covered this, I apologize. I just missed that part.

I agree with the notion of as designed Studies 2 and 3 shouldn't proceed. I agree with the fact that we can redesign or we can recommend redesigning Studies 2 and 3 to be something that's more in line with what needs to be researched. But we also have to be realistic that there's expectations within the community about the work that we are doing. So if we change the scope of our work here, we need to not only let the community know that the scope of work is changing but also let them know how that will impact other processes going forward.

The reason I say this, which is not a surprise, is that many in the Subsequent Procedures PDP have been stating over and over, "Let's defer some of our conclusions to the work of the NCAP," but that was based on an understanding that the NCAP would be doing Studies 2 and 3 as designed. If those studies are changing, that view of Subsequent Procedures may ... I don't know if it will change. That's the kind of thing that we need to have those discussions. But all I do know is that there are others in the community too that keep saying we have to delay and delay the next round until we actually figure out the true – how horrible name collision has been. I think what these have shown us, both Study 1 as well as the previous research that's been done, I think – this is my opinion – that there's nothing, no evidence to be seen that there's harm enough to stop everything else from going to move forward. That could change subject to additional research if we find something, but I think we need to set the right expectations with the community and not just

---

give ammunition for some to say we have to put everything on hold until we figure this issue out. Thanks.

JIM GALVIN:

Thanks very much for that, Jeff. You're right about us being careful about changing scope. I guess my own view of this to offer this up as a partial response here, I don't think we're changing the scope of Study 2 and Study 3. We're changing the questions that we might ask of the data. Is that changing scope? I don't know. I guess that's sort of a matter of opinion. It didn't occur to me, honestly, until you brought it up that that might be viewed as changing scope. But we'll certainly have to keep that in mind as is common in ICANN, it's important that the community be aware of those kinds of things if such a thing happens. It has impacts and consequences all over the place. So thank you for that. Greg, please go ahead.

GREG SHATAN:

Thanks. Can you hear me?

JIM GALVIN:

Yeah. Please. Just fine.

GREG SHATAN:

Great. Thanks. I guess first go beyond my initial reaction to the report, especially with the added ending of basically it sounded like we were all being told to fold our tents and go home, and I understand subsequent correspondence kind of clarified that wasn't quite the case, but it does

---

feel like we've kind of reached an artificial point here I think that the idea that there's no more work to be done just seems off.

I think with regards to the specific point of changing scope, in retrospect, the idea that Study 1 would have no effect on the particulars of Study 2 and 3 was probably the wrong idea back then rather than now that it's a problematic idea that we would indeed look at that and say that both the passage of time and what we in fact have learned through Study 1 and the work of this group today would [intonate] that we've actually could refine what was expected. I don't think that's necessarily a change in scope. Change in scope would really be going beyond the remit into new concepts and topics. Regardless I think if we change the questions, redesign the questions, we can't just do that amongst ourselves and go on our merry way. We need to let people know.

If this were a GNSO working group, which is not obviously, we go back to GNSO Council and say, "Hey, we want to change this. This doesn't work now. We realized that after three years of work that this next task no longer makes sense or we've run into something." We get approval on that through the process pretty quickly and go on our way. If we ask for a massive change in scope, like we've discovered a whole new problem set that we need to work on, that would cause a human cry as it should. I think if we did that too, would there be a human cry? I don't think we're out there looking for new problem sets. I think we're out there realizing that the answer to the judgment which may or may not be entirely correct, but 2 and 3 as designed are not necessary, the answer is then what is necessary to satisfy the larger goals that 2 and 3, and really this group and this work had? Given that 2 and 3 were

---

supposed to be done anyway or were likely to be done, it doesn't seem like that should be considered a delay in the process. Indeed I would hope that given what we've learned in what we've done, we could possibly be more efficient with 2 and 3 than we would have been if we just marched in the straight line that was set, whatever it was, a year or two ago. Thanks.

JIM GALVIN:

Thank you, Greg. I like very much what you said. I want to pull two things out in particular and then give you a chance to if you want to come back into the queue and comment on them. But thank you for also agreeing that changing what we ask in the analysis is not necessarily changing scope, but you correctly pointed out that even if we are going to make any changes at all, even if it is just a change in questions and such, it's appropriate to make sure that we find a way to make that visible to the community. People should know if we're changing something that was in an original document, in ICANN typical GNSO terms, charters often have very specific questions that have to be answered by different working groups and such. In our spirit of cooperation and being open and transparent, if we were going to change those things, we should make that visible, the only thing that's there to the community.

The other thing that you just said which is useful and I want to call out the folks that Matt Larson put into the chat room there, he was quoting out of the report the actual recommendation that Karen has produced there. In fairness, again getting back to the by design stuff that Danny had said originally, and I think that you said it here too, if it's not really

---

properly designed at the moment then let's take the time to figure out what it is that needs to be there and maybe that's the right thing to do. Presumably, that question would get asked. It's also fair to say maybe the Board will say, "Okay, that's not the right thing," and then they'll reach out to SSAC and say, "Please tell me what you need to be different before we decide to fund this." So that would be a work product that we would have to produce regardless. So maybe that's the approach we should take now is producing that work. I just wanted to call that out and let's continue the discussion here. Jeff Schmidt, go ahead, please.

JEFF SCHMIDT:

Hey, everybody. Obviously, I sent a note in support of the Karen draft. This is kind of dovetailing off of Greg's comment a second ago. There's some question about purpose. I feel like this project, we're sort of on a random walk with respect to collisions because I'm not sure there's a clear objective, and that's what I liked about Karen's paper. So let me ask some questions of the group here.

In the latter half of Karen's paper, I don't know, the last two-thirds around line 1400 or so, she has three bullet items. The vast majority of new TLDs delegated since '14 haven't been the subject of any collision reports to ICANN. For every one report, there were eight delegated blah, blah, blah. And during the three-year period, '17 to '19, there was only one report to ICANN. That and then the bar chart or the line chart above that obviously is pretty visually stunning as well.

---

Let's ask ourselves, do we think that there's a problem? If this isn't okay, if we can add 1200 strings to the root five or six years ago and this is the outcome – the outcome is we had a spike controlled interruption did what it was supposed to do, people were aware, there was a media campaign, ICANN wrote advice to people, to administrators, and the curve went down because of all that just like it was supposed to by the way. Since then, since the second half of '17 or so, give or take, we've continued to delegate stuff. I don't know exactly what the numbers are but we've delegated a lot and nothing is being reported.

So I just want to ask and it's not really intended to be a rhetorical question, but are we saying that that was not okay, that we failed somehow? In the period from 2012 to now, did we fail? Is this an unacceptable outcome? If yes, that's fine, but then we have to say very, very specifically where we failed and what an acceptable outcome would be? For every one report in the second half of '14, there were eight TLDs delegated. Does that need to be 1 to 10 or 1 to 20 or 1 to 50? Is zero collisions really the only way we're going to declare success? Until we feel comfortable that we can eradicate DNS namespace collisions entirely, we're unwilling to accept some background noise, that is our definition of success. All those are fine but I think we need to be clear. Otherwise, I feel like we're kind of waddling in the wind here a little bit on what we're trying to do. Thank you.

JIM GALVIN:

Thanks, Jeff. Let me respond to you in the following way. I'm pretty sure I know what Warren is going to say to you in part of your response. But I want to offer up the following as the role of this working group and the



---

way that I characterize the Board questions. Our job is not about the deciding on the success or failure of controlled interruption per se, and it's not about declaring whether the last round was successful or unsuccessful per se.

The Board ran into a problem and what they're thinking about is the fact that in the future, they're going to have to deal with this issue. SSAC has certainly made it clear from its point of view that name collisions are here to stay, they're not going away. They will always be present. So part of your response is, is there background noise? Yes, they will always be there.

What we need to produce for the Board, what needs to come out of our nine questions is guidance for them into the future about how to be concerned about the fact that name collisions will exist whatever the volume is, whether it's 1 in 8 or 1 in 50 or 1 in 1000. The fact that they will exist, how is the Board going to resolve the fact that they exist? Is it okay to simply delegate and not worry about it? Or what are the questions that the Board might ask itself so that it might decide, "Gee, this shouldn't be delegated." So in essence, it's about how does the Board into the future identify the high volume situations like .corp, .home, and .mail so that it can then make a decision about whether or not those should actually not be delegated or not into the future? How does it do that? That's how our job here is coming up with those guidelines for the Board to use to make that decision. Any questions that it can ask itself could get there.

---

JEFF SCHMIDT: I think that's fine but there's a necessary prerequisite there that we've decided that what we did in the 2012 round is somehow not okay. There's some option to say, "Hey, what we did in 2012 was fine. We found the high-risk strings. We cause them to not be delegated and everything else that we did is fine." We can either say that that's the case or proceed down a path where we're implying that what we did in 2012 is not okay and we have to do something different. Thank you.

JIM GALVIN: Right. That might be the conclusion that we come to but it's important to look at everything and think about this thoughtfully and come to that conclusion.

I'm looking at the time here. Okay, let's get to the queue. Jeff Neuman, go ahead, please.

JEFF NEUMAN: Yeah, thanks. I agree wholeheartedly with my other Jeff co-colleague. I think this is all about what we are implying with our studies. Right now the whole fact that this NCAP group was created and the fact that the Board ask these questions has created an implication that we really still have a problem that we need to solve. We don't have evidence yet to support that view. That's not to say that the evidence may not be out. It's not to say that there's no evidence out there, it's just to say we haven't found it. But by implication, the work that we do and the more work we say we need to do, the more of an implication we're saying is that there's still a problem in the way that it was done, the last time is not appropriate or is not the best way forward.

---

I agree with Jim. You're absolutely right. At the end of the day, the Board – most of them are not technologists, many of them are academic or business, they will have a few simple questions. One, are there any strings that in the future we could never delegate because they present such a high collision risk? Two, are there any ways in which we can test other strings if we don't know all of the high-risk ones? Is there a test we could apply that would give guidance during an evaluation that would say, "Hey, this one presents a risk"? And the third is, can we mitigate that risk of the string so as to allow those delegations or can we not so as not to allow those delegations? It's pretty simple. It's a simple question, not a simple answer, sorry. Very difficult to answer.

I think that's where we need to keep the focus at least for the Board. If we want to do other research in how new technologies impact other things that are related but on a tangential relationship then we need to be honest with the community and not just visible as to what we're doing but why we're doing it and what kind of message we're trying to send. Thanks.

JIM GALVIN:

Thanks, Jeff. Very good point. I do agree with your questions. As you said, simple questions, not necessarily simple answers. I'm going to jump to the queue here rather than responding more directly, just to get some other voices in. Warren, go ahead, please.

---

WARREN KUMARI:

Thank you. As Jim said, probably many people know what I'm going to be saying here. But I'm still very uncomfortable with relying on the number of reports that ICANN received about issues with collisions, because we have a lot of evidence, much of it is anecdotal, but we have a bunch of evidence that they were collisions that were not reported. When people had collisions, what they said was, "This isn't working. This broke my system. I'm having the following set of issues." But we didn't see those being reported to ICANN for a bunch of reasons. In many cases, people didn't know where. In many cases, they didn't want to expose this externally. And I suspect in a large number of cases, they just figured, "This isn't going to do any good. Why should I go through the effort of filing a report saying I was using the string internally and now my system is broken?"

The way that controlled interruption was set up – I think I might be stepping on some toes here – but it seemed as though it was specifically set up to not gather any sort of useful evidence. There were a number of proposals on how we could collect evidence so that we would know if this is working, if this situation is getting better, the magnitude of the problem, but that simply wasn't specified in the final controlled interruption design. So we don't have the evidence now on what the magnitude of the issue was. So this continuing going back to – "There weren't that many issues reported, therefore, they weren't issues," I think is a false analogy. I think that if another round is done, the way that we did controlled interruptions was not sufficient and we do need to collect a bunch more evidence. Whether that's something where instead of it being 127.0.5353, it's directed at a public address so that an external party can collect data. Or whether whoever is doing the

---

controlled interruptions is required to self-report a bunch of data. Those sorts of things I think are needed. "We don't know if there's an issue because we didn't collect data," I think is not the way we should be doing this. But I think that's my standard rant.

JIM GALVIN:

Thanks very much, Warren. The quote that you've used before which I like, which I know has been used for other things is, "Absence of evidence is not evidence of absence," and it's important to keep that in mind. I think that part of our role in this discussion group is in fact to examine what effect that particular issue has on the whole system. Maybe controlled interruption really was perfectly fine and we should continue [following] that. But I do think that it's appropriate to actually look at that and look at the data and at least draw whatever conclusions we can for whatever fact and evidence are available. They may be incomplete just as whatever decisions were made back in 2012 or made based on best evidence available at the time, and it's appropriate to do that again and prepare for the future. There are a couple of other gaps I want to highlight but, Danny, go ahead, please.

DANNY MCPHERSON:

Yeah. Can you hear me, Jim?

JIM GALVIN:

Yes, just fine. Thanks.

---

DANNY MCPHERSON:

I was literally going to follow on verbatim to what Warren just said. I agree with that. I think that using reports to ICANN on a model that wasn't necessarily peer-reviewed or otherwise is certainly included feedback from the community and looking at the reports to ICANN and that is not the way to comprehensively measure the efficacy of collision mitigation frameworks across the board. It might be the efficacy of measuring controlled interruption and who reports to ICANN once they do a Google search and find out what they need to do. Controlled interruption – there were some classes of attacks that I believe they did a fine job notifying people. But there's a whole lot of classes of attacks that it didn't. All the DNS service discovery protocols and everybody that's going to be open to a man-in-the-middle attack because of one of those protocols, it isn't going to see that controlled interruption failure. Something is just not going to work or is going to work and they're going to be compromised. Who's going to report those compromises to ICANN? Nobody is. That's the real threat, the real risk besides breaking stuff.

Jeff Schmidt asked the question, he says do we think there's a problem? I don't know, Jeff. My response would be, do we think that corp.com is a problem? Then the question for that would be, why do we think it's a problem? A big part measurably of why we believe it's a problem is because .corp is a problem, and you can look at .corp now and see tens and thousands of strings of service discovery string at the root for queries that occurring that are the result of potential collisions that could be exploited by an attacker. If .corp looks like that, what other TLDs look like that? If other TLDs look like that, like .mail, .home, and .corp, yeah, they're high risk strings. But before we enter a new round,

---

shouldn't we make sure for the benefit of the applicants and the benefit of the clients that are going to be using those TLDs that if we can prevent those attacks in the first place and prevent people from expanding years and years of work waiting on those strings to be delegated with their business model to say, "Here are the [rails] for this," not just, "We can test it with controlled interruption for 90 days before it's delegated because that's not what I subscribed to."

The other thing I would say is that if we don't believe collisions are a problem then the paper, Jeff, that you just shared with this group said that during this time when people are working from home, the number of queries at the root have went up for TLDs where IT administrators have misconfigured things or ignored the risk from those. That's not all .corp. As a matter of fact, there are a lot of other strings there where we see this behavior for. So I think to ignore this and just say we controlled interruption and declare a success because ICANN didn't get reports is a very, very myopic perspective. It's all I got to say. Thank you.

JIM GALVIN:

Thanks, Danny. I agree. Moving on to the queue. Jeff Neuman, go ahead, please.

JEFF NEUMAN:

Thanks. I agree with Danny's questions. I think the questions ... We should stop doing all this stuff around the edges and just get directly to Danny's questions. Are there strings out there and how do we figure it out? All this other stuff of compiling all the resources and all that other stuff and gap analysis, new technologies, I mean, let's figure out the

---

precise things that are needed to get to the question that Danny asked. Because we can't account for everything in the future, if there's some sort of test that we can do, after someone applies for a string, where we could make a similar type conclusion, let's redesign it that way. But we've sort of spent a year or more around those issues and not getting to it. My fear is that the way this community works, because it's extremely permission-based, is that people are going to want permission to launch the next round, rather than just doing it if no evidence has been found. If we can set out a straight scope and a timeline for answering those questions and then that will be communicated to the inner community at ICANN, let's do it. That's just the right way to do it.

JIM GALVIN:

Thanks, Jeff Neuman. Continuing the queue, Jeff Schmidt, go ahead, please.

JEFF SCHMIDT:

I agree, Danny. Remember, when I said, "Let's look at what we did" – "we" collectively – "in the 2012 round," one of the things that we did that was a great success was find and hold the delegation of .corp, .home, and .mail. I've referred to those over the years as unicorns. What I've said is it is essential that we find those unicorns. We probably won't be able to determine them in advance. We'll probably have to look at a string by string analysis. We need to find those unicorns just like we did before and prevent delegation of those. Everything else is just noise.



---

Sure, we can make improvements, the paper that I just wrote about corporate IT departments. You know what? Collisions are a bad thing, they should be fixed. The most harmful collisions don't happen in new gTLDs. Just by numbers and by size, the most harmful collisions happen in comm, and that's fine. We should help people understand this issue and fix it.

I'm not disagreeing but the issue is again the "for what purpose." For what purpose do we exist? I think we know how to delegate strings to the root at this point. That doesn't mean the collisions are fine and we all just whistle past the graveyard. But I think we know how to delegate strings to the root right now. On that very specific question, we would have to say that what we did in 2012 was somehow inadequate with the specific use case of delegating strings to the root in order to then move on and decide we have to do something different. Thank you.

JIM GALVIN:

Thanks, Jeff. I see Danny and Warren in the queue. I want to put a line into the queue at that point. I want to be sensitive to the time here. I want to try to wrap up a little bit at that point. I still don't believe I have any answer to my question about the next steps. Let's get the comments on the table and then try to summarize them. Danny McPherson, go ahead, please.

DANNY MCPHERSON:

Jeff, I appreciate that response, actually. I think that the crux of one the unicorns, if I look at our root query data and look at what exhibits things like large volumes, in my opinion of DNS service discovery, web traffic,

---

types of protocols, that kind of thing, there's a long list of things there. Most of its stuff people don't want, but also a lot of its stuff like .corp, .home, and .mail. I agree with that, if we can identify those things and how we get there.

Jeff Neuman asked question that was like, "How do we test this?" and I agree with that. I think that's the crux of where we're at and you sort of look at the context.

One of the things that VeriSign listed in their references to the working group was academic research, peer-reviewed paper that we tested and employed internally at VeriSign and looked at exactly that. We said, "Which of these strings exhibit these behaviors?" That's how we wanted to identify these as high-risk, high-volume strings. I think that part of the problem here is nobody has a clear offering for .corp, .home, and .mail and nobody knows how to identify those unicorns in the next round. I think that's the way. I think you can't test those with controlled interruption. You need a mitigation framework like was published in that paper, maybe not very recently ago, but in 2016 and 2017 and I think it's a useful contribution to the working group into the community, something like that. Quite frankly, it's really no different than what [Interall] and Lyman did and what Duane Wessels did in 2009, and then what Jazz and Verisign and others did in our work along the way. Saying, "Hey, here's protocols that seem to be trying to learn something from the DNS." If we can figure that out and move on, I think that's it.

As for the studies, again I agree with the ... I said all along that I wasn't sure that Study 2 or Study 3 in our current form are going to be useful

---

that we just need to employ the frameworks and do the work for the string as they come up. I think that's kind of probably what needs to happen as well.

JIM GALVIN: Thank you, Danny. Warren, go ahead, please.

WARREN KUMARI: Thanks. Warren Kumari responding to Jeff Neuman. I mean, I think what we did in 2012 was inadequate because we didn't actually collect the data. The way that the controlled interruption was designed, it specifically was not collecting data that would have told us information on whether collisions were actually causing problems. I think that, at the very minimum, the next time we do this, we need to collect data so we understand when controlled interruptions stop, who's being affected, and if they're actually suffering issue or if it's just queries that weren't useful anyway. That's basically my summary.

JIM GALVIN: Thank you, Warren. Well, secondly, Jeff, I closed the queue, but I'll tell you what, I'll give you the last word here. Let's just speak for a minute and then I want to try to wrap this up.

JEFF NEUMAN: Thanks. Well, maybe I could just go to AOB because it's related. I think it was inadequate. Also, you're bringing the technical perspective and I'll bring the commercial perspective. It's inadequate to allow people to

---

apply for strings, only to find out years later that it's deemed to be ineligible because of these things that we could have known prior to starting the round. So I think it's just another way to look at it and another way we need to approach the problem too, because this does have impacts from a commercial perspective. I think the .corp, .home, and .mail problem is an interesting one that we do need to solve but I think that's going to be a very different issue than the going forward issue when there are delegations. Remember, these inadequate delegations are not just happening with new gTLDs, they're happening with new ccTLDs that are delegated every few months or several a year. There are new IDN ccTLDs that are delegated that just continue to happen without this kind of scrutiny, to the extent that we could study these things on a rolling basis. In fact, if we can figure out a way to get certain data out of the strings that are yet to be delegated even before we start a new round.

JIM GALVIN:

Okay. Thanks, Jeff. Let me try to pull this together here.

JEFF SCHMIDT:

Jim, could I have 10 seconds? Sorry. Since the honeypotting thing is an issue here. I'm so sorry. Jeff Schmidt. I totally hear Warren. This is a discussion we've been having for years. I would say look at our Phase 2 paper, we talked about this issue. Taking something that is not sending traffic out on the Internet and making a change to cause computers to send traffic out on the Internet of unknown confidentiality, when it was not being done before, is a gigantic liability issue. We talked to lawyers

---

about it. It's a big issue. So honeypotting, yes, it gets more data but it's not a magic panacea. We didn't choose not to do it because we don't like data. Thank you.

JIM GALVIN:

Right. Thanks, Jeff. Since I let you say that, I'll let Warren say what he wants to say. I'm pretty sure I know it's coming. But try to get through this quick. Go ahead, Warren.

WARREN KUMARI:

Thank you. I'll go through it quickly. Delegating the name does make the victim in this case start sending data on the Internet. But potentially, they're sending it to a malicious party, instead of somebody who's set up specifically to run the honeypot presumably in a safe way. I think it's much, much, much, much, much more dangerous to say we're going to let these names be delegated and then allow miscreants to go up and register the name versus setting up a honeypot where, yes, there's legal liability but the data is being sent to somebody who should be trusted. As soon as you delegate the TLD, people can go and register malicious names and now the data is being sent directly to them. It seems like a much, much worse outcome from a security standpoint and potentially from a legal standpoint. Going, "We don't know if this is going to do bad things. We're not willing to collect the data ourselves. But hopefully attackers won't do it." Feels not good.

---

JIM GALVIN:

Okay. I want to cut off this discussion. I'm sorry, Danny, about whether or not we're going to talk about the honeypotting or not honeypotting. I want to leave that out, try to summarize this here. I think the fact that we've been having a really good sort of detailed discussion about some of the analysis that has to happen, I want to remind us what the role of this discussion group is. We've talked about things like predictability of applying for a name. Well, part of our role here is to provide some guidance to the Board on how to decide if a name that exhibit collisions should be delegated or not. Part of that analysis for us is to actually say whether or not controlled interruption is a good thing that helps that analysis. Then you package up data that you get to the Board to review and make its own personal assessment about these things.

I think Rod Rasmussen had a comment early on in the chat room that I wanted to call out here. This is not about whether or not controlled interruption was appropriate or not but it is fair to ask the question, "Do we have enough data to decide whether or not controlled interruption actually served a purpose?" That was one of our technical analysis gaps in our brief was the observation that, "Well, as limited as it is, we definitely should do a thorough review of what we know about controlled interruption and the things that were identified so that we can form whatever conclusion we can about that and make whatever assessments." Let's see. I think that this worked with this discussion group to do. That's one thing that I wanted to say.

The second thing is we need to keep in mind that there are some things that exist today that did not exist in 2012 when they decided to run around and do what they did. That is in today's world, we also have to think about whether or not the system can be gained. Whatever set of

---

rules we create for assessing, whether or not a string should be eligible for delegation or be required to be considered a collision string, to use the Board terminology, and thus not delegate it. One of the extra pieces that we have to think about is the criteria that's used to make that assessment and if it gained or not. We have to give some consideration to that. That was not a data point back in 2012, because they were dealing with data that was there before and collision strings were not on top of everybody's mind at that time. Again, there's still work to be done.

I think the third point here for me in all of this is that we've been making the comment that you can analyze this to death, you can collect all the data you want, but when you're done and when is it over, how do you do all of that. I think my answer to that is one of the things we pointed out in our technical analysis, our technical gap brief was the idea that the infrastructure has changed and we should in fact expect that the innovation will result in the infrastructure changing. So let's take a look at what has happened given that the infrastructure has changed between 2012 and now. Let's see if there's anything discernible in that for us in looking at metrics that we might propose the Board use to decide whether or not to delegate. And can we thus provide any advice about how to review this process at regular intervals to confirm that it's still meeting their needs? That's part of our consideration, too. And that is certainly something which was not an issue in 2012.

There's two activities and predictability that we have to consider that is beyond just whether or not a string is delegatable or not. It's also about the future of this process. I hope I didn't confuse that too much. The

---

---

critical question I want to get an answer to that we didn't get today was whether or not this discussion group feels like it's important for it to say something in addition to the work product one and its final report and the recommendation that Karen's making. I don't believe we've talked about around this a lot and I don't think we've come to a hard conclusion. I want to continue this discussion next week. I do think it's important to have this discussion. I think the Admin Committee will take as an action item here to try to construct a statement of position that we will put out to the list so that folks can actually have something to throw stones at or agree to. I find the best way to work in the sort of volunteer open groups is give people some black and white to look at and that tends to crystallize things. It either moves us forward or it makes it very clear what the path forward is. We'll take that as an action in order to get us to next week.

That gets us to the top of the hour. I know that Jeff had an important any other business. Jeff, let me let you go ahead and do that. Folks, could you could hang on for just a moment for Jeff's comment.

JEFF NEUMAN:

Thanks. Tomorrow, we have a call. I believe it's 20:00 UTC we being Subsequent Procedures PDP Working Group. We've invited Jim and Patrick and Matt Thomas and Matt Larson. Actually, some of them are coming, but if you want to attend that meeting, just to see the work that we're doing within that group, you should come. Please send an e-mail to Julie Hedlund or myself if you want to go. We'll get you the call-in details. Thanks.



JIM GALVIN: Okay. Thanks very much, Jeff. Any other business from anyone else? I apologize for being quick and short. I'm not seeing any hands. I think that's an old hand, Danny.

Thanks, everyone. I appreciate your flexibility here at the top of the hour. We'll see you all next week. We're adjourned.

KIMBERLY CARLSON: Thank you. Bye

**[END OF TRANSCRIPTION]**