

KIM CARLSON: Thank you. Welcome to today's NCAP Discussion Group Call on December 11th at 20:00 UTC. In the interest of time, there will be no roll call. Attendance will be taken according to those on Zoom. No formal apologies have been received.

As a reminder, all calls are recorded and transcribed and the recording transcripts will be posted on the public wiki. Again, to avoid any background noise while others are speaking, please remember to mute your phones and microphones. And with that, I'll turn the call over to you, Jim and Patrik. Thank you.

JAMES GALVIN: Great. Thank you very much, Kim. Patrik Faltstrom, one of our co-chairs here is also on the call with us today. Probably worth noting that we still have an interest in seeking a third co-chair. So if anyone is interested in that, we would encourage you please to step up and let us know.

Let's do the usual ICANN kind of thing here when you have formal working groups like this and just ask if anyone has an update to their SOI that they would like to announce and let everyone know has changed.

And not hearing anything, folks can just keep that as a reminder that we do actually have to track that and you should pay careful attention to that if you ever have changes in your employment status, usually is what makes that matter the most. But if anything relevant comes up,

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.

this is your opportunity to always announce that at the beginning of these meetings.

We do have one new member, Justine Chew, who has filled out an SOW and put it in. I don't see that she's on the call with us today. Just so that folks note, she is also now a part of our discussion group here.

KIM CARLSON: Jim?

JAMES GALVIN: Yes.

KIM CARLSON: Just for the record, Justine was originally a member. She asked to be an observer and she just recently asked to become a member again.

JAMES GALVIN: Oh, okay. I see. You wrote that there. Justine updated her SOW. Okay. So thank you for that.

Okay. With that, as we have said sort of ambiguously, I suppose, a little bit here, certainly in Montreal and even on our list a bit here, our goal here for this meeting is to jump in and look at our, essentially, ten questions or ten identified questions, if you will, that come out of two resolutions from the Board and to walk through those with two particular goals in mind. One is to make sure that we all understand. We have an understanding amongst ourselves as to what we believe that

question means in case there is any discussion or questions about it, we want to come to some consensus about what actually is being asked of us.

And then number two, our goal is to identify what we think we need to know in order to respond to that question or topic. I want to be careful to point out that we're not trying to answer these questions. That's actually the role of Study 2. That's the next phase of this project. That's not this phase. So if we take a look at these and we understand what we need to know in order to answer that question, the goal with that information is to use that to feed into Karen, our consultant and contractor here, is to feed into her work.

She has two responsibilities, two sort of over-arching objectives, responsibilities in what she's doing in Study 1. One, of course, is collecting whatever she can find with related to name collisions and documenting that and cataloguing that, so looking for any kind of published work that talks about name collisions in any way and then creating a summary of that, a bibliography, if you will, of everything that she can find. And ideally, if folks here have any suggestions, comments about that, you should add to that list. She did send us a beginning list here. She started that process by letting us know what she's working with. Folks should please take a look at that and look to see if there are things you want to add or if you have any questions or comments about anything you see on the list. Let's have that discussion.

But the second part of what she's supposed to do is to look to see if she can identify any technical gaps that would contribute to the decision of whether or not to go forward with this project, meaning Studies 2 and

Studies 3, as currently laid out in the overall project plan. And our goal in looking through these ten things is to think about what data we think we need in order to respond to those questions. And with that, I would expect that the data that we think are gaps, that would feed into and influence her summary report of technical gaps that she sees because if she can't see that that data is available in anything in her bibliography, then that does become a gap that does need to be looked at and considered as we make a decision about whether to go forward with Studies 2 and Study 3.

So let me pause there, see if anyone has any questions or comments about that. Does that make sense? And, of course, Patrik, or if you want to add anything, please do go ahead. And I'm not seeing any hands at the moment.

Okay. So I think then the plan is to just start at the top here and go through these. Although, I think I'm not going to start at the top. I'm going to jump to the second one on the list, which is line three, which is now in the display. Everyone should have seen the document.

We have an external drive here in which we've been collecting our questions. As we create these kinds of documents that are useful to the community, we'll publish versions of them onto the wiki so that folks in the community can see them. For now, everyone who is here should have access to that document. If you don't, make a note here in the chat room and we'll see if we can deal with that in the background. But otherwise, Kim will keep it displayed up here on the screen.

With respect to the definition of name collision, if folks are keeping up and I'm going to presume that everyone has been, there has been a definition of name collision that was established and it was developed over the course of the first couple of quarters, up through May, I believe, of this year. It was published and put out for public comment and all of those public comments were reviewed and so we have an established and set definition for name collision at this time. It was drawn out of what SSAC had previously written and updated based on the fact that time has passed and we know some new things.

So we're not going to go through that anymore at this time. However, it is fair to acknowledge for the group here that although we considered that particular step done for right now, if you have new information you want to bring to the table, then you should certainly send that to the mailing list and then we'll consider that and consider whether or not we're going to open up that definition. But for now we consider the definition of name collision to be a done activity and a done task for this discussion group at this time.

So that leaves us with nine things up here. Understand what they are. Let me pause there for a moment. Any questions on that step? And then we'll go back to the top and start. Okay, not seeing any hands.

So there were two resolutions from the Board. One resolution is actually Row 2 here in this spreadsheet. The second resolution is the other nine items which are actually labeled inside the item, one through nine. The first one is the Board asking us to speak about .CORP, .HOME, and .MAIL and, in particular, speak about the risks posed if those names should be

delegated and any mitigation that might be used with respect to any identified risks that we can put there.

So a question for us at this point is do we believe that we have enough data that we know enough to be able to speak to that in Study 2 when we get to actually doing our analysis or is there data that we would like to see?

I have one suggestion to put on the table that I will put out there. I think that in the case of this, and probably this relates also to some of these other questions, one of the data sets that is, hopefully, we can find a way to have access to in some way even if maybe we can just get some analysis done on the data, even if we don't get the data itself. But there are a number of very large publicly accessible resolvers that, in addition to the route server operators, which is data which was available just prior to the last round of new gTLDs, we now have come into existence a number of publicly accessible resolvers and I think that it would be interesting to be able to ask of those resolvers, questions similar to what we asked of the route server operators back in the day.

We don't know for sure that we'll be able to get access to that data or get answers to those questions, but I do think that that's new data and I do think that it's important for us to be able to ask those questions about .CORP, .HOME, and .MAIL of that new data set if we can so that we can also consider that in our response here.

So that would be my one suggestion about data, a gap in data that I would like to be able to see. Any comments about that? And Warren, you have your hand up. Go ahead, please.

WARREN KUMARI:

Just mentioning that I think a number of the large resolver operators have pointed out the fact that they have very strong privacy protections around that sort of data and it's sort of possible that they might be able to say something along the lines of, "We still see a really large number of these," but getting anything much more detailed than that, I suspect is going to be really hard.

I'm also not entirely sure what getting something like that would actually say. If you see that there's five of these a second or 5 million of these a second, as with many other things with name collisions, you don't actually know what the impact is going to be because you don't know what the implications of getting anything else back for the name are.

Also following up from that, the ICANN Board had a number of times strongly implied that these names were never going to be delegated and also, as far as I know, suggested that the original applicants retract or withdraw their applications. I'm not quite sure what would happen to people who have done so, if any have. If these names were then suddenly put forward as things that could possibly be delegated, but that seems that would be opening a fairly terrifying can of worms. Luckily, that's not my problem though.

JAMES GALVIN:

So thank you for that, Warren. Let me respond to your last comment first. Yeah, I think I want to keep this discussion group focused on the technical issues and technical concerns. So the business issues of

contracts that may or may not, or applications that may or may not have been withdrawn is really not of concern to us. And yes, the interpretation of a previous Board resolution on whether or not these names would be delegated is also not something which is in front of us to address or be concerned about.

On the first half of what you said though, you're right. I am sensitive to the fact that there might be privacy concerns, certainly with accessing the data. We had accounted for that when we originally wrote this project proposal and so there are things to be considered there, things to be worked out. We may be limited in what kinds of answers or responses we can get to questions. We may be limited in what questions will even reasonably be answered by those resolvers and their data sets. However, we at least have to acknowledge that there is an opportunity there and we'll have to address privacy concerns when they come up and we're ready to do that. So that is important.

I'm actually trying to figure out how in this spreadsheet here. Oh, it did work that time. Okay. I'm trying to figure out how to get a carriage return in there so I could put this up there.

Okay, I think that covers it. So that's a fair thing to note. Rod, you have your hand up. Go ahead, please.

ROD RASMUSSEN:

Yeah, just thinking a little bit about this. I think I agree that this a great source, potentially a great source of data. It depends on what people are willing to share. The more particular and narrow questions are, like the issue around .HOME, .CORP and .MAIL may be easier for some

people to answer than some of the more generic questions like, “What are the queries that you’re seeing that are undefined that would be potential collisions?” and things like that where you start getting into things that may be more sensitive.

I think that figuring out an inventory that people to query on, to have discussions with about what kind of data they might be able to share and having good questions to propose to them is probably how I’d want to address this, both for this question and other questions that are further down the list. Thanks.

JAMES GALVIN:

Yes. Thanks, Rod. I agree with you. I believe that part of Study 2, the analysis phase, really is about us crafting questions that we want to ask of data that we believe will inform our analysis and our ability to make recommendations.

I didn’t really want to do that right now. I don’t think that we need to do that now. For me, and I’m open certainly for reactions to this, I believe just acknowledging that there are data sets and that we know that we’re going to prepare questions to ask those data sets is appropriate for right now. I don’t think we need to identify the questions that we would ask at this point in time. That would be something that we would do as part of Phase 2 and moving into that.

Let’s see if anyone wants to react to that. Rod, go ahead, please.

ROD RASMUSSEN:

Yeah, and I largely agree with you. I think an exercise in, at least, thinking about potential questions leads you to thinking about what kind of data sense to pursue. But I think that we know, at this point, that those recursive resolvers can answer lots of questions. It doesn't really help you with that. There might be some other things that thinking about some potential questions could jar as far as other data to look for. I'm not sure what that is, but that's just something, food for thought. I'm not saying we should run off and do a bunch of questions right now. It's just something to think about. Thanks.

JAMES GALVIN:

So fair enough. I think what I will do here as part of the discussion is take note that maybe, let's go down to the rest of these because I think one of the things that you said before is we may find that it's helpful for us to understand the kinds of questions we might ask or at least get a few questions in mind that we might want to ask of data. And let's see how often, if that comes up again as we go through some of these. Maybe we will spend some time in the future in another meeting thinking and drafting some potential questions that we might want to ask because you're right. Maybe that'll help us to think about data sets that we need in addition to just the resolver one, which I sort of jumped out at here but maybe there are other things and that will peak our interest in that.

And okay, so with that, Warren, you had a hand up but you took it down. So I'm guessing if you leave it down, you don't want to speak.

Let me just make a note here right in the spreadsheet. Okay. Maybe we need to find a better way to do this in the future instead of holding this up.

Okay, so we're going to skip number one and we're going to... So to close off that discussion, first of all, before we move to the next one, a critical question here for the team is do we understand the question that's been asked, and do we have anything else to add here that we think would be helpful to us as we consider how to answer this question. And if not, we'll move on to the next one. But Warren, you have your hand up. Go ahead, please.

WARREN KUMARI:

Yeah. Just mentioning that if there are questions for [inaudible] resolver operators, the sooner those can be produced, probably the better because it's going to take a while for the privacy implications to be discussed. But also, when the questions are drafted, if they can be as narrow as possible, things like, "How many things do you see matching .HOME, .CORP, or .MAIL?" are potentially things that could be answered. So just notes for when questions are eventually discussed. But it could be a long time to get an answer through the privacy side. That's all.

JAMES GALVIN:

Sage advice. Thank you for that. Okay, captured that point too.

Okay, with that, I am going to move on down to Row 4 and Item 2 on the list here, which is the role that negative answers play in the

operation of existing end systems. In fact, arguably, this is really, from an operational point of view, this is one of the things that caused name collisions to really stand out. .CORP, .HOME, and .MAIL, probably most notably .CORP, but in any case, there was software which was dependent on the route servers actually negatively responding to the existence of the name. And they depended on that, on that particular behavior. And so it's useful for us to speak to this now.

Some of this has actually been described and explained previously for those who remember the site finder, wild card discussion that happened back in the early 2000s. There's an SSAC document where we talked a great deal about how changing negative responses, and there's quite a lengthy document which talks about negative answers and what that means and the expectation on the part of existing software for that behavior. And this is one of the reasons. It's really one of the primary motivation reason why the wild card responses, wild card records are not permitted in a TLD, at least not gTLDs anyway is because of expected behavior about negative answers.

So this would be about us probably, we would have to consider whether there's more to say here, but I think what we'd want to do, primarily is go back to that site finder report and adapt that material into the context of name collisions in particular. But it strikes me that most of the information that we need for dealing with this is going to be in that report. At least, that's just my personal view about that. I'm certainly open for other comments and suggestions from folks. And Steve, go ahead please.

STEVE CROCKER:

Thanks, Jim. Given this in comparison what Karen wrote about studies to take into consideration, the site finder related materials she sent, at least in her documentation, she hasn't decided whether to include that in her comprehensive analysis. So maybe given this, we may want to, the working group may consider whether to ask her explicitly to review the site finder materials where negative answers returned changed. Thanks.

JAMES GALVIN:

And I think, Steve, that's an excellent suggestion. I have not double-checked, but I'm assuming that you're actually, yeah, you're pointing out that the site finder stuff is not on her list. But we should take a note here that Karen should include the SSAC site finder report. I apologize while I take a note here. Okay, that's my answer to that one and the note that I would take there for that in the spreadsheet. Any other comments, suggestions from anyone? Okay.

With that, we'll move on to the next item on the list, number three. This one actually gets a little bit interesting in my mind, at least, anyway. So they're asking us to talk about, the Board is asking us to talk about the harm to users if collision strings are delegated. So what is it that actually can happen, whether accidentally or purposefully exploited? What are the consequences, the downstream consequences of a name collision?

And this again, I guess there's probably a thought exercise that needs to go on here. There is some data available about all of this and I suspect that what Karen is collecting is going to provide some information in response to this. I also expect that we will have spent some time in a bit

of a thought exercise about what we know from existing literature and also as the technical experts that we are, and of course, we now have several years of experience with all of this, I think that it will be an interesting tabletop thought exercise to consider this. It's not immediately obvious to me what other data to ask for. We'll certainly have access to anything that's been reported and we can look at that and see if that piques our interest any further on considering options. But those are the two things that occur to me, incidents that have already been reported. There's a handful of them that ICANN has recorded and then, of course, whatever the literature currently says that we've learned over the past half-dozen years of new gTLDs being put out there.

Interested in comments, questions, suggestions from others while I take a note here. Looking for a hand. All right, so those are my notes for that. Interested in other thoughts. Can anyone think of any particular data that might be useful to us or helpful to us in answering this question? Are there other questions that you might ask of something else that would be helpful? Are there any related areas that might be helpful in us thinking about this? And then, of course, the general question, do folks believe that they understand the question being asked? Okay.

Sorry that you had to listen to my voice for an hour, but happy to let anyone else jump in who wants to along the way here. So let's move down to the next one, Item 4 here.

So possible courses of action that might mitigate harm. This is actually what Study 3, or the third phase of this project overall, is about. The purpose of Study 3 is to look at potential mitigation methods and to

consider them and, in fact, actually to see if we can act on them, if we can exercise them. That's the word I want to use.

There is this idea of trying to build out a test system, a test harness, and potentially be able to try some of these mitigation methods. So I think that for here, as we identify possible, as we identify sources of name collisions as it comes out of the literature and any bit of experience that we've had over the past six years, we also need to take a look at, we need to review the incidents that had happened, look at what might have happened. We also do need to do a review of—what's it called?—controlled interruption. I apologize. Controlled interruption also needs to be reviewed here because that's the current mitigation action that was actually chosen in the last round. So that's an important thing to review here, and then also to consider if other courses might be available.

What's important here is thinking about the future, what needs to be able to happen in turning this into predictability for the delegation of new strings is, is it possible to understand how to evaluate new mitigation methods that might be proposed in the future. And I think that's an important part of considering new factors and courses of action is also providing for the opportunity for reviewing things that might come up that we don't know.

Any thoughts or reactions from anyone?

PATRIK FALTSTROM:

Yes, we do have controlled interruption that we should look into, like was said in the chat, to try to evaluate how effective that was. It is also

the case that if we go back in the SSAC discussions and also the SSAC proposals, there were some other proposals that have already been designed and put on the table at the time of selection of controlled interruption. So I think there is material around the discussion that led to controlled interruption that could give some ideas on other mitigation methods. And I think it's fair to put some effort into look into why those were not selected in those days and reevaluate them.

JAMES GALVIN:

Yes, thank you for that. Let's see. Where would the documentation be? We have controlled interruption. We have the SSAC report. I know that she is also looking at the [jazz] report which had its suggestions. Where is the documentation for other mitigation things that worked? I'm not aware if anybody did anything else.

PATRIK FALTSTROM:

No, Jim. I think that is what we are talking about. My only point was that it's not only controlled interruption that is hidden in that material. There are also many other proposals as well, or many, a few others. That was my point.

JAMES GALVIN:

Okay, thanks. Okay. And Warren, you have your hand up. Go ahead, please.

WARREN KUMARI: Yeah. I think that Verisign, for one, had written some document on, or done some research on mitigations and their potential utility or not. Unfortunately, I don't have a link in front of me but if there is anyone from Verisign on the call, maybe they can poke at it or maybe just reaching out to them would be useful. Sorry I don't have anything more specific but I know I have seen a few other documents.

JAMES GALVIN: No, that's helpful. Thank you, Warren. Matthew's there. Matthew, go ahead please.

MATTHEW THOMAS: Yeah. I believe, Warren, that you're talking about one of the documents that we posted at the workshop prize name collisions. I'll dig it up there but it should be included in the literature for review that the contractors included, I believe.

JAMES GALVIN: That's good. Thank you, Matt. I appreciate that. In fact, I'd appreciate if you would just make sure that that showed up on the list that Karen was sending around. And otherwise, send it to the mailing list. Send the link to it again at the mailing list so that we have that. In fact, I'm going to just say that you're going to look for that. Okay.

STEVE CROCKER: Jim, the other thing is in Appendix A of [SAG] 66, the SSAC lists four, sorry three alternative approaches and SSAC did analysis on the pros

and cons of those. So that's [SAG] 66, which I think is included in Karen's documentation. Also [SAG] 62 where the SSAC lists two broad categories of mitigation. So those are in [SAG] 62. Thanks.

JAMES GALVIN:

Thank you for those references, Steve. So that's good. So we certainly will have some material to consider there and to look at. Okay. Anything else from anyone? And let me see. I'm reading Justine's comment in the chatroom here.

I'll add those words that you're offering there. I think my immediate reaction, Justine, is that the last three things we've listed there are all sort of sub-parts of evaluating alternatives to controlled interruption. But we'll just list it out separately right now. We don't have to be too structured about this at the moment. It's most important just to make sure we don't miss anything and we collect everything that's important to us. Okay, so thank you for that.

Let me just catch up with the chatroom here from Anne. Do we have any way of knowing what happened with end users after the 90 days of controlled interruption?

Well, I think the answer to that, and I'm looking for anyone here to correct me on this, but we only know what people reported and so if no one actually said anything, then the assumption here is that we don't know because remember what they got in response was a pointer to a very special IP address that if anybody was paying attention, they would go look up on Google. They would Google that and then they'd get

some information and go do something, and in theory, if they reported it, then we would know something.

So arguably, it was successful if you didn't get a report and that kind of thing. Warren, do you want to add something? Go ahead, please.

WARREN KUMARI:

Yes. Just want to mention that there were a number, there were some things which people posted on things like various forums. Stock exchange had a bunch of reports and things which were things like, ".DEV is suddenly not working for me," or ".PROD is not working." I think it's really unfortunate that the controlled interruption system was set up the way it did and that there wasn't a webpage provided where people could go and log this sort of information, or at least, have the registries have to report how many hits they were getting for names that resulted in a controlled interruption response.

So there aren't any formal things, but if you look around on Hacker Overflow, sorry, Stack Overflow, and a bunch of other ones, there are a bunch of reports. I'm not quite sure where I saw a spreadsheet, but somebody had started creating one and it got to at least 100 entries of anecdotes of "I had a system that was working, .PROD was delegated, and suddenly a bunch of my systems stopped working." Things of that ilk.

I do think that if another system like this is set up, it would be really useful for much better data collection to be done, potentially with some sort of way to point people at a place to report stuff. As you said, you get a special IP address and maybe people notice they will go look it up

on Google and then know something but that doesn't actually provide any information back onto the occurrence of the issue.

JAMES GALVIN:

So thank you for that, Warren. You're right. I do understand that there wasn't any data collection at the time. There was, in fact, a conscious choice to not really provide any data collection outside of reports associated with controlled interruption, unfortunately. Well, or fortunately, depending on your point of view, I suppose. The path that was chosen was really based on a great deal of emphasis on respecting privacy considerations.

So that certainly is something that we could take under advisement. As part of our consideration and analysis, maybe we can say more about that and make better recommendations going forward. Or another round or any future rounds that come with that, maybe we need to balance the needs of name collision a little differently against the potential for privacy concerns. That's certainly a topic that we can have some discussion about.

I do want to come back to this Hacker Overflow, Stacker Overflow stuff. Is there a way to track that down and get that and make sure that Karen has it so that she can catalogue it and include it? Any chance, Warren, that you could look for those references? Warren, go ahead, please.

WARREN KUMARI:

So I'll see if I can find a spreadsheet, but simply searching for 12705353 finds a bunch, or just searching for controlled interruption, etc. As they

say, they are just anecdotes though and the plural of anecdote is not evidence. But there is certainly a bunch of them. And if anyone else can find or remember the spreadsheet, that would be helpful because I don't actually know where, what I'd even search for to try and find it.

I think that there... Now that I'm talking about it, I think that there was a presentation at a DNSSEC [OR], or maybe DNS symposium which had also just a list of some other anecdotal responses.

JAMES GALVIN:

Yeah. It's good that we're having this discussion because I'm sitting here thinking to myself we need to make sure that this kind of stuff gets brought to Karen's attention and gets tracked down and included. Now I put the three links that you just put in the chat up here. We need to make sure that those get to Karen. I can do that from the links that are here.

Your comment about the DNSSEC [OR] presentations, I remember some pointers to presentations earlier when we were collecting potential documents and stuff. But I don't know in particular if these made it there. Let me make, let's see, one comment here. And let me insert another link here. I'm just copying all of your great links out here. Warren, you're putting them here. Let's get them out and in here so at least there's a record of them and then we'll need to make sure they get on the mailing list here too so that they can get to Karen and she can catalogue this stuff.

Okay. Keep those links coming. All of this is good. Any other comments or questions about courses of action and any mitigation that might have

occurred to make sure we get all of that and capture as much of that as we can?

I think that if we get good answers to Item 4, Row 6 there, the possible courses of action and the things that people did, then we're going to have to, our objective in our analysis phase is going to be to actually look at all those things and review them and consider what factors affect the success of those courses of action. Again, keeping in mind that our goal here is about general guidance about how to evaluate future mitigation methods and, of course, we have to consider specific advice about controlled interruption and what else might be possible.

So I think that this question, Item 5 here, in evaluating the potential success or failure of courses of action, it's really about getting good answers, a good set of data from Item 4. And with that, then we can review all of that and do a thorough analysis of what happened there and the circumstances. I think that's our best bet there. Plus, of course, any other thought exercise that we might invent for ourselves here about potential mitigation methods. Any comments or reactions to that?

Okay. Not seeing any hands, let's move on to the next item here. Doing a quick time check, but I think we have plenty of time still. Item 6, potential residual risks of delegating collision strings even after taking actions to mitigate harm.

So I think that one of the things that we need to do, the way that I see this question, is it's important for us to consider the entire lifecycle here. So a name, there are potential collision issues prior to a name

being delegated that we might be able to identify and discuss. There are things which we'll know about as a result of the pre-delegation period, whatever that happens to be, if it's controlled interruption or something else, there will be mitigation actions that will probably be allowed and they will be executed.

But there's also the longtail side of all of that. Even if you're doing all of the right things, it doesn't mean that you can take care of everything. Maybe there will always be... There will be some kind of longtail and I read this question as us having to give due consideration to the longtail problem of collision strings. If a string, which is potentially in this collision category is delegated, there are risks, there is mitigation, but what does longtail look like? What can we say and respond to about that? And Steve, you have your hand up. Go ahead, please.

STEVE CROCKER: Can you hear me now?

JAMES GALVIN: Yes.

STEVE CROCKER: Good. With respect to this question of consequences or mitigation, taking a page from other things that are happening quite visibly around us, suppose that a careful evaluation suggests that it's okay to delegate one of these strings and that the proper precautions have been taken in mitigation [inaudible]. And then there's a sharp negative reaction from the public or from some portions of the community that ICANN is selling

out or has biased it in favor of the financial motives of the various parties.

So that I would suggest that the evaluation needs to take into account two things. One is the actual harms as best as can be evaluated, and then the reputational harm that may be somewhat separate from any objective evaluation.

JAMES GALVIN: Okay, thank you for that. I'm trying to capture that. Actual harms and reputational harms.

STEVE CROCKER: Yes.

JAMES GALVIN: Good point.

STEVE CROCKER: And the point I'm making is that you could have an evaluation that suggests there's tolerable or negligible actual harm and you may still incur reputational.

JAMES GALVIN: And I think that's a good point. People have varying views of ICANN's role in all of this space. But setting that aside and looking at it objectively, reputational harm is not necessarily a technical issue. But I

do think it's important to point out the fact that that is something that ICANN does need to worry about or we should at least mention it. And we'll see how far we can get into it given that our primary focus is going to be one of technical issues. But for completeness, it would be appropriate to make sure that we at least mention this in their presence and where they fit into the overall picture of things.

Okay. Anything else with respect to Item 6? Okay. Steve, I'm going to assume that's an old hand and move on to Item 7 here.

Suggested criteria for determining whether an undelegated string should be considered a string that manifest name collisions in place and the category of collision string. Now this is the part that gets, I think, really the most challenging for us. And this, at least for me, when I think about this list of questions, this becomes the question which I suspect will garner a lot of discussion and I expect some varied points of view on this. I interpret this question to suggest that what we need to do is provide guidance to the Board and the community at-large as to what kinds of information, what kinds of, well, criteria—I was trying not to use the word criteria but I can't think of any other way to say it—so what kinds of criteria are available that could be used when trying to decide whether or not a string is a collision string or not.

Everything that we're going to be doing up to this point is looking at the kinds of things that cause collision strings to exist, what do we know about how that can be managed or not managed. And from that, we have to develop some guidelines for how to evaluate strings and how to determine whether or not it's going to be in this category of collision string. This doesn't necessarily mean it will never be delegated. The

point here at this point is simply, how do we decide whether something is likely to be a collision string or not, what kinds of data can easily be suggested to go get.

So even for us here, we're looking at what do we want to look for in root server data. What might be interesting to see in publicly accessible resolver data? Is there a way for us to suggest that these kinds of things should be, they should be asked for and collected as part of preparing a package for the Board when it's going to evaluate a given string? That is sort of our model here.

So that's something for us to think about and we're going to build this as we go along. I don't have anything in particular to suggest for this. I think it's something that will develop as we go, and as we look at everything else that we've got here. So Steve, you have your hand up. Go ahead, please.

STEVE CROCKER:

Yes, I do. Thank you. So I'm considering the question fresh and trying to imagine somebody else reading this. The most obvious thing is that the string shows up in queries to the root with some frequency. You can raise the question as to what that frequency ought to be, but I want to go to a slightly different question which is what else could we possibly imagine would be a trigger for deciding that there may be a collision issue? I can't think of anything and unless we can suggest something, then treating this question as just very broad strikes me as... There's a word I'm groping for, but disingenuous is too strong. But I mean it seems to me that the [inaudible] that is of issue is how often it occurs

and then one can ask the question, why is that and what does that suggest is the issue underneath that? But unless somebody can come up with something that runs that question, I would recommend that we either narrow the question or specifically point to the frequency of existing non-resolvable occurrences.

JAMES GALVIN:

I'm just trying to figure out a way to capture what you're saying. So okay. I'm just... I'm not trying to be too overly solution setting. What I'm trying to do here at the moment is capture the kinds of questions that we need to be thinking about to drive us towards how to respond to this. What I'm hearing you say is we should be careful not to broaden this question but to think about it fairly narrowly, lean more on the narrow side than on the broad side when we think about this issue and I'm trying to capture your points here on that. What else besides root server presence might be interesting to consider? But it occurs to me that as part of our analysis of all the prior things, let's look at how often did name collisions occur, why did it occur, is there any way to evaluate the frequency of those occurrences to drive us to any recommendation?

STEVE CROCKER:

Yeah. I think you've got it basically, Jim. What I was suggesting was kind of two parts to this. One was speaking to the obvious, which is to [inaudible] queries to the root. What is there to say about frequency and any patterns of that? And then the second is, is that all there is? And you've got both of those in there and you can weave those into, rather directly, rephrasing Question 7, to speak to how to make use of

the frequency of occurrence and also raising the question, any other basis for selecting a string that might be considered?

JAMES GALVIN:

Okay, thank you. I do want to do a time check. Take note of we're almost at the top of the hour. But we have some hands, so let's run through the hands here. Warren, go ahead, please.

WARREN KUMARI:

Thank you. See, I think I'm just responding to Steve's thing. I think that there are a bunch of other things which cause, or potentially cause, issues. For example, there's all of the, there's the IANA special use domain names registry which has things like .INVALID, .ONION, .TEST. Those don't necessarily show up at the root all that often, but they are sources of collision and if they were delegated, bad stuff would happen.

There are a bunch of similar things, which I think things like HTTP. I suspect that if that were delegated, there would be some things that would end very poorly. So I don't think that it is just queries that show up at the root. I think there is a larger set of things that should be considered for a name to be delegated. And I'm sure there are a bunch more. Those are just off the top of my head.

JAMES GALVIN:

No, thank you. That at least gives us a path to be thinking about as we go down through these. I might generalize HTTP to say any sort of protocol label might be interesting, but that's worthy of some

discussion. This is enough to keep us on that path when it comes up later.

Rod, you have a hand up. Go ahead, please.

ROD RASMUSSEN: Yeah, I'm just adding on to that. I think the distribution of the queries, and also what we're looking at as far as mitigatability, if that's a word, of that. If you have something that's very narrowly focused but having a five frequency of queries to the root, that might be a solvable problem versus something that's a low level but it's spread everywhere. So I think that that's a consideration. Thanks.

JAMES GALVIN: Say that again, Rod. Distribution of queries...

ROD RASMUSSEN: Yeah, they're coming from a lot of places or there may be particular kinds of places that you want to make sure. If everything's coming from .MILL that's a potential collision, you might want to think about delegating something, for example, the .MILL [ES] space, so to speak.

JAMES GALVIN: Okay. I think I captured that. Just take a look at that. Thank you. Matthew, you have your hand up. Go ahead, please.

MATTHEW THOMAS: So beyond just frequency and source distributions, one of the things I'd like to comment on, number seven and going back to number one, is what questions we might want to ask will likely also be influenced by the things that have happened in the DNS ecosystem as a whole since the last time we did this name collision analysis. We've had technologies that QNAME Minimization, NXDomain cut. Even we have things like DOH and DOT going on there now that have likely influenced our observation capabilities at the various layers in this DNS hierarchy. And so that end, what do we need to consider in those ramifications in terms of how our future analysis and measurements should quantitatively and qualitatively assess those compared to our base lines that we use longitudinally.

JAMES GALVIN: That's a really good point. I'm going to capture that here. How does changes in the DNS ecosystem affect our analysis? I'm going to use DOH as a good example as one to start with. Oh, and yeah, the QNAME Minimization, two good ones. Okay, thank you for that. Warren, go ahead, please.

WARREN KUMARI: Thank you. I think what... I'll try and reinterpret what Rod said just because I'm not sure if it came across completely. I think what he was trying to say is one site sending 50,000 queries a second is different to 50,000 sites sending one query a second. And Rod, please correct me if that's not at all what you meant.

ROD RASMUSSEN: That is what I meant.

WARREN KUMARI: Something else that's... I'm following on from what Matt said. The RC 7706 and 7706 [inaudible] stuff is also going to be really important for our visibility. There are a substantial number of resolvers which currently aren't talking to the roots at all for determining NXDomain and using aggressive NSEC as well means that the root simply isn't going to see these. So the fact that queries aren't hitting the root for .CORP doesn't mean that they're not actually occurring. It just means that they're being stopped by local root or hyper local root, whatever you want to call it, and aggressive NSEC. And that's going to, I think, significantly [inaudible] account, and unfortunately, I don't think there is a way to tell how much that's occurring.

JAMES GALVIN: Okay. I'm aware that we're over time here and... But all of these are good things. Okay.

So one quick logistical thing to ask. Thanks very much for this discussion. This has been great. We got down to a really good one here. I think what we'll do next time we meet is give some people some time to think about this. We'll definitely allow for a quick additional review of these first 14 rows here, if people have anything they want to add, and then we'll jump in to talking about eight and nine.

Quick question, do we want to meet next week or just move into the new year? And I'd like to frame that question in the following way, and

folks should be able to use their Zoom spot here. You should be able to. I believe you can, under the... Or maybe just raise your hands. I'm looking and it doesn't seem to want to work. Yeah. Are you available? There it is. Everybody got a poll there at the moment. Are you available to meet next Wednesday? On yes, or no, or unsure. We'd like to sort of get a quick look at that and consider and see what we get with that. We can do a quick poll here if folks can take a moment, please, to quickly click on that and then we'll give Kim here a chance to tell us what she sees and decide quickly here if we're going to meet next week or just wait until the 1st of the year. Whenever you're ready, Kim.

KIM CARLSON:

We're still only at 80%, 88% of people voting. But I'll go ahead and close it and share the results.

JAMES GALVIN:

Okay. So we have, at least of those on the call today, we have a good percentage that'll be here next week. So let's say that we're going to meet next week and then we'll also continue into the new year. After January 6th, whatever Wednesday, I think that's January 8th if I've got Monday date right. But that first full week of January after the first of the year and we'll commit to weekly meetings, and we'll take up meetings as we need to. So we'll go from there.

And I see the comment from Justine in the chat. We'll just take that for the moment.

So thanks, everyone. I apologize for running over. I appreciate your willingness to stick with it here. We'll see you all next week. We're adjourned.

UNIDENTIFIED FEMALE: Thank you.

UNIDENTIFIED MALE: All right. Thanks, everybody. Bye-bye.

[END OF TRANSCRIPTION]