
KIM CARLSON: Thank you and welcome to today's NCAP discussion group call on February 5th at 20:00 UTC. In the interest of time, there will be no roll call. Attendance will be taken by those on Zoom. Kathy and I will update the Wiki with the attendance as quickly as possible and put the link in chat.

We have a few apologies: Chris Roosenraad, Eric Osterweil, Ram Mohan, Matt Larson, Steve Sheng, and Roy Arends. All calls are recorded and transcribed. The recording transcripts will be published on the public Wiki. As a reminder, to avoid background noise while others are speaking, please mute your phones and microphones. With that, I'll turn the call over to you, Jim. Thanks.

JAMES GALVIN: Thanks, Kim. My co-chair, Patrik Fältström, is actually with us today so you have us both here today. We'll just tag team as we go through the day, here, through this particular meeting. Next up is to ask if anybody has any changes to their SOIs to announce. I'm not seeing any hands go up, so moving right along. Kim, I should have asked you a bit before if we had any new members. I'm guessing not, or you would have mentioned it, but let's just double-check.

KIM CARLSON: Yeah. No new members as of today.

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.

JAMES GALVIN:

Okay. All right. Thanks very much. A bit of formality with normal ICANN process working groups, here. I apologize. Last week, we canceled the meeting relatively last minute. Our sincere apologies to folks, especially those in the Asia Pacific region who probably didn't know the meeting was canceled until they woke up in the morning extra early and tried to join. It just had got overtaken by events, there, and neither Patrik nor I were able to be present. I appreciate your flexibility and tolerance of all of that.

We're going to jump right in, here, to our one overarching action. We had the draft document distributed. A couple of versions have been out there. Folks should have had plenty of time to take a look at that, offer any comments that they have. There hasn't been too much discussion on the list. I don't know if Karen has gotten any other direct feedback that she might have taken on board. Ideally, folks would have sent it to the list, not to Karen. But independent of that, we now have our action. There were two particular actions in study one.

One, of course, was the catalog of existing work, as much as we could find. Anything that might be related to name collisions. Karen has certainly done an outstanding job of that. Thanks to everyone who has been feeding references to her. The other action is to provide some advice on whether or not studies two and studies three should be funded and should go forward, and whether or not we need those in order to do the analysis that we have to do in this discussion group, anyway.

Part of feeding into that, of course, is taking our spreadsheets, which hopefully folks have opened in front of them. We have asked about this before. We've done this a couple of calls ago. Our objective, here ...

We've had some discussion about what these ten questions from the board actually mean.

Now, what we need to do, the process that we were hoping to go through, here, is to provide references back into this draft work product, here, for each of these questions and, with that, take the discussion, the backreference, and create a question that we need to ask of data or, "What is our analysis question, or questions, that we need to ask ourselves so that we can answer these particular questions from the board?"

What do we need to do? What information do we need? How are we going to get it? Is it available from this work product? If not, is there something that we can do to get it? That's what we need to build out, here, so that that can be included in this work product.

That will be used to decide whether or not to fund the additional work that needs to be done. Although, in any case, we still have to do whatever analysis we can do with what we have with respect to these questions. Let me pause there for a moment and give Patrik an opportunity. See if he wants to add to any of that, suggest this differently, and, of course, see if anyone has any questions. Put your hand up if all that makes sense. Patrik?

PATRIK FÄLTSTRÖM:

I don't have anything to add, thank you.

JAMES GALVIN:

Okay. With that, probably the easiest thing to do is start at the top and work your way down. I'm really hoping to get some discussion from people in this. I really don't want this to be a limited show, here. Ideally, folks have come prepared and are ready to do this and contribute to this work. That's why we're here.

I'm really looking for folks to put up a hand and have something to say in offer of all of this. If folks need more time because they weren't really prepared for this, that's fine. We'll simply pick this up next week. I'm not going to sit here and just vamp on all of these things myself.

Going down the board questions, here, going to row four, it's labeled as question two. Since question one is actually already done, we have a definition of name collision, one that we're working with for now. At best, this group might decide to revisit that definition. But we'll revisit that as part of our analysis when we move into doing that phase of study two, if we think we need to revisit it based on what we've learned in all of this, here. But for now, we're just going to leave that as a done action and focus on question two, here.

We had a discussion point, here, of Karen including the Site Finder reports, which she has done and she has put in the report. An important question for us here is, do we have what we need in order to answer this question? Or can anyone think of any additional question they might want to ask about negative answers and understanding that experience as part of our analysis that we'll do in Phase 2? We had said to ourselves before that the Site Finder report likely has everything that we need. I personally tend to agree with that point of view but I'm open to suggestions from anyone else. Oh, I see a hand. Warren, please go ahead.

WARREN KUMARI: I think I've mentioned this before but we should also note that ... I don't know if this is the right place to note it but we should also note the fact that aggressive NSEC caching means that some of the negative answers aren't actually hitting the root. They're not really returned from the root, for these. They are returned from local resolver. I'm not quite sure how we word that, though. Did that make any sense or did I just mumble in a way that nobody understood?

JAMES GALVIN: No, it made sense to me. I'm going to write some words here in the discussion thing. You said NSEC caching affects ... This is new information as compared to what was covered in Site Finder. I don't believe this was covered in Site Finder, which I think is your point. So NSEC caching affects the visibility of NXDOMAIN at the root?

WARREN KUMARI: Yes.

JAMES GALVIN: Which is the only place that we looked for NXDOMAIN, in root data.

WARREN KUMARI: Yeah. So both the aggressive NSEC caching and the QNAME minimization both make it so that the negative answers aren't actually seen at the root.

And so, knowing how this will affect end-users is potentially a little trickier.

JAMES GALVIN: So the proposed gap here is that we will need to add discussions about NSEC caching because that's not covered by Site Finder.

WARREN KUMARI: Yep.

JAMES GALVIN: I'm just adding those words over here in column F. There we go. And let's change the format to wrap that.

WARREN KUMARI: I guess NSEC caching and, maybe, QNAME minimization.

JAMES GALVIN: Okay, sure. Okay. All right. So that's a gap in something that we'll have to do in our analysis but I don't think that affects the work product. What's interesting is that might be even existing root server data that we could get access to. Even this year's DITL data, or at least the most recent DITL data. I don't know. Maybe we can find an additional root server operator who might actually do some looking at some data for us, even if they don't give us the data. We should look to get some new information. I would think that that's possible. Matt, you have your hand up. Go ahead, please.

MATT LARSON: Yeah. I was just going to piggy-back off of Warren's comment and say that we might want to even broaden that from NSEC and QNAME minimization to general technologies that would impair NXD observations, things like local root or hyper-local root and other technology changes like ... What is it? NXDOMAIN cut, where it's similar to QNAME. And in terms of the analysis at root, Verisign has invested heavily in retaining and storing longitudinal data of A and J. So if we frame certain questions, I can definitely poke at the data.

JAMES GALVIN: That's excellent, Matt. I appreciate that very much. I'm glad to hear that. I'm saying, here, "Other technology changes that have affected visibility of names at root servers." I'm just going to leave it at that for the moment. I think we will ... I'll add that same thing for here for the gap. Public technology change discussion. Okay. That'll be some additional work that we'll have to do. Okay. Excellent. Thank you. Anything else from folks?

There would be some data to query, there but it might not be ... We'll have to see. We'll have to get a collection of data queries that we might need to execute and see if that warrants a ... If we can just do it amongst ourselves because we have a few people who are willing to make queries against data that they have or if we actually need to go down the path of looking for a contractor who will then try to collect all of this from a variety of sources. We'll have to make that assessment later after we have a list of these things. Okay.

Next one down, unless anyone has anything to add? I'm just going to keep moving along, here. Oh, I'll add a comment about local root back-up there, too, when I get there. So looking at question three, "The harm to existing users that may occur, collision strings were to be delegated." We do have some notations about incident reports in our draft report. We should pull out that reference.

There's probably not a lot here. I mean, there are some references to some work that has been done. We'll need to pull those out. This is a case, at least for me, the way I'm thinking about it, this might be a thought exercise for us. Maybe we can extrapolate from what little we do know about what we've seen and what's been documented.

But I expect that we'll have to spend some time and a bit of a thought exercise about what happens, especially the way infrastructure-level protocols work when collision strings are delegated, and what happens. We'll have to give some thought to that.

I'm open for other suggestions at the moment but I don't know that there's much else we can do. We have a record of what we know about and the analysis is likely to be us spending some time in a thought exercise. Hoping for other suggestions, though. Okay. I feel like I'm doing all the talking, here. Oh, Matt. Please go ahead.

MATTHEW THOMAS:

That sounds like a good plan. The only other thing I wanted to note is a few minutes before the meeting I did send out an e-mail with some potential questions, thoughts, or avenues that we might want to explore

specifically around this with regards to the reports in section four and five. We might want to take a look at that if we get a chance.

JAMES GALVIN:

Yes. As input to this discussion, yes, of course. I'm going to open that up. Okay. I'm opening the message up and I'm looking at it. I don't know if other people are, too. This is comments about the document directly. At the moment, we're going through the spreadsheet, here, so I'm not sure how to jump into these questions unless, Matt, you think that your questions are relevant to something we're doing in the spreadsheet.

I think, for the moment, I'm going to suggest that we keep focusing on the spreadsheet here. But I'm open, if people want to jump into this, if they've had a chance or if they have thoughts. If people could look at the e-mail, if you think you have discussion points, we can jump into that.

MATTHEW THOMAS:

I'm fine with that.

JAMES GALVIN:

Okay. And I'm sure that folks should look into what they want to say about that. Okay. Item four: possible courses of action that might mitigate harm. This, of course, is a direct reference to the fact that, for now, there is no real guidance that exists in this. This is an opportunity for us to consider things that we can do. This, too, might be a thought exercise except for what we can find in any incidents that have been documented.

Keep in mind that study three is all about it's possible trying to test, or at least considering various mitigation options as a thought exercise, and then commenting on them and talking about their efficacy, and wonders what circumstances they might be appropriate or might not be appropriate.

And in fact, we should provide some thoughts about the controlled interruption, how it has been working, what it does, and whether or not it's the right thing. Certainly, that's what's being done right now or what was done in the previous round of new gTLDs.

And the question is, going forward, do we allow that as a test to see if they exist? Do we allow for ...? One of our guidance options here has to be, if name collisions exist does it make sense for there to be a mitigation proposal that the board can take on board and decide that, "Okay, with this mitigation strategy it's okay to delegate this, even with the existence of collisions." That's the space that we have to examine with respect to this question.

And so, we need to think about whether or not it's possible to mitigate the harm and what might make sense to do and what might not. A valid conclusion, from our point of view, is that name collisions can't be allowed at all, that there's no way to have a strategy that allows them to be made safe enough for delegation.

That's a question that we have to answer and that's the space that we have to explore. This, again, is probably mostly a thought exercise on our part, extrapolating from what we know about collisions that have occurred.

ROD RASMUSSEN: Jim, I've got my hand up.

JAMES GALVIN: Yes. Please, go ahead, Rod, while I'm typing notes.

ROD RASMUSSEN: Yeah. On that, I agree. I think we know this but it might not [hard] to write it down, here. That would probably be some sort of per-string test capability, right? We're not talking about a blanket solution/policy. It's more of a ... I'd almost call it a framework for determining if your mitigation is available for that particular issue and then whether those mitigations are viable for reducing or eliminating the risk. We're going to have different things cropping up for lots of different reasons, some of which may be trivial or easy to take care of and others may not be.

JAMES GALVIN: That's a good point. Thanks for that, Rod. I added a second bullet in the proposed gap column there about, "Might be dependent on string type and quotes," meaning that, yes, we probably have to define what that means, too. But in other words, perhaps there's a mitigation framework that would be helpful for evaluating these things. That would be the space that we're going to get into, here, in answering this question.

ROD RASMUSSEN: I'm thinking of classes of collision types and mitigation types, something along those lines.

JAMES GALVIN: Okay. Classes of mitigation types.

ROD RASMUSSEN: Yeah. Well, collision types and mitigation types. Having things like software issues, configuration issues, or things like that. And then, mitigations can be of various flavors, as well. It may be something like patches, redeployment, or things like ... There are different categories, is what I'm thinking there. That might be getting into the weeds a little bit but it's, I think, good to have that framing so when we are thinking about these things we're not trying to come up with a one-size-fits-all solution.

JAMES GALVIN: That makes sense to me. I changed my words a bit there to say it might be dependent on classes of collision types and mitigation types. That is, perhaps, as a mitigation framework that would be helpful. That's what we want to develop. But in any case, I think for the most part this one is really a thought exercise extrapolated from what we happen to know.

We have a variety of sources, here, and we will have to review each of these carefully and see what we can pull out of them that might be useful. We'll see what else comes of that, okay? Any other comments from anyone?

Not seeing any hands so we'll move onto number five, factors that affect potential success of the courses of actions to mitigate harm. This is probably just a part of item four, and I think we said that in our discussion already, here. Careful review of item four is what is needed here.

To me, this is an ordinary risk/security analysis kind of thing that has to be done with any of the elements that we're talking about in item four. That's the way I think about this. I'm certainly open for other suggestions, here, on how people think we might approach it. But I think it's an element of item four and our review of that is to consider the risks.

It's really a risk analysis on whatever courses of action that we happen to come up with in item four. Okay. Unless we have any additional comments from anyone, I'm just going to keep plowing on through, here. All right.

Question six: potential residual risks of delegating collision strings, even after taking actions to mitigate them. In the spirit of what we wrote down there for item five, this is also, probably ... Yeah, the discussion comment that we had here, before. It's just a deeper risk analysis consideration of the long-tail problem of the risks. I don't really think there's anything else to add except what we had here in our original discussion.

Again, this is probably primarily a thought exercise on our part to walk through some of these things carefully, whatever we can identify and list. It's probably worth pointing out the obvious, here, in some sense, which is, what about things that we don't know? We don't know what we don't know. I think that's an important thing for us to be keeping in mind as we continue with this work, even when we're doing some of this analysis.

We always have to account for, what happens if something comes up that we hadn't thought about and ends up not being documented? What's the default or fallback action here when other things happen that were not previously considered? This is that question in my mind, for us to think about, really, "What can we do long-term?" and how to respond to those kinds of things. What advice would we give to the board and ICANN in general in that situation, to consider? Okay. I'm not seeing any hands or anything in the chat so we'll move onto question seven.

"Suggested criteria for determining whether an undelegated string should be considered a string that manifests name collisions." This where we get into what's probably the most interesting part of our analysis. This is the part which is the most visible to the community. In here, what was obviously stated in our discussion up here, right upfront, is that the way that the NXDOMAIN queries looked at the root are the most obvious consideration.

.Corp, .home, and .mail, because their obvious, significant number of queries, number of NXDOMAINs, that come to the root, really do put them in a special category as compared to all the rest of the collision strings that had previously been identified in DITL data.

An important consideration here is, to make it concrete, is there a number where collisions under this number that are seen at the root mean it's okay to delegate, and above this number you can't delegate? My guess is there's no such number. I'm not suggesting that we need to come to that number. We need to think about, what are the characteristics of that number? What is the guide that we can provide here so that, in any given situation—and every situation should be

unique—what are the questions to ask so that we can evaluate that? We've already had all of that discussion.

This, again, represents a thought exercise. I don't know that there's any additional data that we might want to get at in asking this question. It occurs to me that it might be useful, in addition to DITL data ... Again, this is where you get into resolver data, kind of thing. We do have a number of large, publicly-available resolvers and there is a significant fraction of Internet traffic that passes through them.

And so, it might be useful to understand what traffic looks like there, when name collisions are present, when collision strings are present, versus what it looks like at root servers. It seems to me that this is another time, another opportunity, for us to understand more of what's really happening. This is an opportunity where being able to look at resolver data and understand the characteristics of collision strings there would be helpful. Warren, you have your hand up. Go ahead, please.

WARREN KUMARI:

I've said this a number of times but I guess I'll just repeat it so it doesn't get lost. For pretty much all of the large resolvers, there are going to be some very sticky privacy considerations around this sort of data. And so, if ICANN or whoever wants to get access to this data, asking sooner rather than later would be good so that lawyers have time to discuss how we deal with this, what data is actually shareable or not, etc.

All of the large resolvers have fairly strong privacy protections in place for user stuff. Having looked at the Google one, for example, there's very

little we can say other than, “Yes, there are some queries that we see sometimes.” Sooner rather than later for that sort of question is good.

JAMES GALVIN:

Thanks for bringing that up, Warren. I do know that that’s an important point and you’re not the only person who has said that. For our purposes, we get to think about what access means. Part of when we wrote the project proposal in the large, we allowed for access to be interpreted in a couple of different ways. I mean, one way, of course, is that ICANN gets a copy of the data and then we can just do whatever we need to over the data, iterate over it as needed to get some analysis.

Another way, of course, is that we get access to the data as in the owner of it just makes it visible so that analysis and the questions can be asked of and by an external third party who is contracted to do that and we iterate as needed.

And then, the third possibility is that we, instead, have to carefully craft questions that we’re looking for. And then, we ask questions. And then, we ask of the resolver people, yourself as a particular example, “Would you ask this question of your data and could you give us a response and tell us a little bit about, as an aggregation, as an overall metric ... Not the details, of course, but just a summary of what does this look like in your particular resolver?” And we will have to sort all of that out.

One last comment. Perhaps this is a more detailed discussion we can have in our face-to-face meeting that we’ll be having at ICANN67. We can dig in on the details and try to think about those questions and get a head start on them, which I think is your point. We should understand what we

want from the data and that's something we ought to do sooner rather than later. Maybe that's one of the thought exercises we should dig into when we have our opportunity to sit face-to-face in Cancún.

Although, I'm certainly open for making that one of our weekly meetings, here, if that seems to be the best way to do it. Your hand went down, Warren, so I'm guessing were okay with that response. I got ahead of you. We need to make a note here to ourselves in this question. "And particularly true for resolver data." Okay. All right. Anything else on question seven?

We're just flying through things here. All right. Question eight: suggested criteria for determining whether a collision string should not be delegated. These questions, seven/eight, are obviously related in the same way that five and six are related. One is the criteria for whether an undelegated string should be considered a string that manifests name collisions.

So it is a name collision string. It is a collision string. How do we make that decision? And then, of course, part of all of that is, should it not be delegated? Or maybe in the future, at some point, there is criteria that will allow it to be delegated. There's the management of this collision string list. Even if it's identified as the collision string, maybe there's a mitigation strategy that will allow you to clear the space for the use of the string.

That's something that we will have to allow for in our framework, here, so that someone can apply for it. It might be there's a process for dealing with mitigation. And then, at some future date ... They can't get it right

away, but at some future time it's off the collision string list and they can delegate it.

All of these two questions, seven and eight, are closely related and we do have to consider, in our framework, who to respond to these things. And what should the board look for in understanding a given string in the context of these two questions? Again, I think this is a bit of a thought exercise.

It might also be that there are questions that we can ask about how collision strings behave that we want to understand – behave in the root server, what we see there, what the data looks like, what they look like at resolvers, collision strings, if we can ... We still have .corp, .home, and .mail that we can probably ask about as ideal candidates, for the moment, to help us get some understanding of the characteristics of how these things behave and maybe anything else that we can come up with along the way.

I don't think there's anything to add to eight except that it's part of item seven. And again, a thought exercise except for we have to come up with questions that might be useful from resolver, and again from root server data.

Okay, that brings us, actually, down to our last one. We're just cranking through, here. "Measures to protect against intentional or unintentional creation of situations which might cause strings to be placed in a collision string category." As we said over here in our discussion ... Well, we didn't actually use this phrase in here.

But one of the things that we considered in the project, overall: this gets to the question of gaming. What can we, once we try to establish a predictable evaluation process ... And that's our ultimate goal, here; try to create some guidance which provides a little bit of predictability or as much predictability as we can get.

I mean, ultimately, some of these decisions become subjective and we're just going to suggest that some data should be gathered for collision strings and then provide some guidance on how to evaluate it. So there's going to be some subjectivity. But at least the process of evaluation will be predictable. There's always the question of, can people game the system? Especially if you're dependent, primarily, on query counts and query numbers.

And we need to speak to that. That's the issue that we're trying to get, here, in this item, question nine. That's what's our responsibility in, again, a thought exercise in how the data can be manipulated. Whatever the questions are that are going to be part of this guidance, can the answers to those questions be manipulated?

We have to do our risk analysis on all of that and provide some guidance about that. Hopefully, with any luck, even providing a way for that to be detected or at least known that it might have been gamed. Rod, you have your hand up. Go ahead, please.

ROD RASMUSSEN:

Yeah. We talked about this a little while back and there are some good notes there in the discussion. It occurs to me a gap that we may be able

to get some answers to are, what kind of instrumentation and visibility could we employ to look for patterns for suspected gaming?

If you think about if somebody was to set up a fairly unsophisticated gaming, which, given cybercrime history, shows that people do simple things because they're lazy, we might be able to find some things. If you take a look at the timings, the traffic that's showing up, is it periodic or does it look artificially created?

There are things you can do to actually investigate that if you have visibility into that. That actually might be one of those areas where you could actually go to, say, the resolver folks who may not want to share data, but say, "Hey, we've got this thing that we're suspicious may be artificially created," and ask some questions about the kind of data they're seeing, something like that.

And I think it is still a thought exercise but understanding how things are instrumented, so not just looking for raw counts of things but timing of counts and things like that where you can look for patterns is important. Thanks.

JAMES GALVIN:

Thank you for that. That's actually excellent. I tried to capture those two points here in some notes. What kinds of instrumentation could be employed to identify attempts to manipulate? And then, of course, perhaps the obvious thing about working with root servers and resolvers. We really do want to understand what could they, and would they, be willing to do as part of this in order to help contribute and be a part of

managing the security, stability, and future name collisions in the DNS. We'll just have to see what that turns into. Thanks for that.

Okay. I'm just going to make one little note here in eight. All right. Well. That was one discussion. I apologize, you've had to listen mostly to my voice for all of this time, here. But I think that we've actually had a pretty good suggestion, here. I think if I try to tie all this, come to some conclusion here ... This is nice. This is good. We probably need to find a way to turn this into something.

I'd like to turn all of these things into something that can be included in the work product that Karen is producing. I'll have to take that back as a consideration, unless somebody has some suggestions to offer at the moment. Patrik and I can take back as a question to consider what to propose here to the group on moving forward with that. That's one.

The second thing that's important to me is I know that the current draft report suggests that there are no additional data sets that would be useful. I do think that we have identified at least a couple of questions that would benefit from some additional understanding of ... "What name collisions look like at resolvers," is the way that I'll phrase it. So there are questions of resolver data that could be asked that would be useful to have in the same way that DITL data was studied.

It just feels like there ought to be some questions we can ask of resolver data and see if we get similar kinds of behavior and characteristics as what we know and understand from the root. And again, as has been pointed out up above, the operation of the DNS in the large, especially at

the top levels, really has changed since the original JAS report was done and the [NRI] report.

Because we do have NSEC caching, sometimes aggressive caching. We've got the local root stuff and QNAME minimization. It really does change what data looks like and what data looks like at different locations. In fact, this is one of the reasons why QNAME minimization in particular makes resolver data more interesting than root server data, because they will certainly be a movement towards root servers having less and less data to provide and to offer. We need to talk through that and think about it, and make some recommendations.

I think that there are some very specific gaps that we need to capture, here, to get into this work-study one report for Karen to include. We need to propose a way to capture that and put that out there. That's certainly interesting.

And anyone who wants to contribute to that effort, would love to have you just type into this document. Feel free to move to the next column over in any one of these things and propose a question that would be interesting to you. As you think about some of these analyses that we have to do in our proposed gap, if you have more details that you want to offer please type it in directly.

Everyone who is a panelist should have direct access to this document and you should be able to go in there yourself, and you should be able to edit it. Please, if you can't, make that known in some way. You can alert myself or Patrik, or Kim Carlson is actually probably a better person to tell. Let her know that you're having trouble accessing and she will work

with you to sort that out so that you can see it and you can make your entries directly. Please add any additional details here that would be helpful.

Now, Anne, I see that you have a question in the chatroom, here. Yes. You're pointing out the SubPro work: "Do our notes take this into account?" Anne, as I recall ... "Possible courses of action." I think we mentioned SubPro in one of our discussion items here. It has certainly been on my mind. But you know, you're right, Anne. Thank you.

SubPro is not actually called out explicitly here in row six. We're scrolling up to row six, Kim. [I had a question four]. I thought that somebody had ... I actually don't remember if SubPro ... I'm a little confused here, myself. I don't remember if the SubPro stuff was included in the work product, now.

Because I know we've had that discussion, there. That's something, Karen, if you're listening, to check, too. There was a note to the mailing list, I think, at some point, which might have come from you. Yeah. And I know that Jeff is commenting that SubPro is on his mind, as well. There was a note on the mailing list about recommendations and SubPro.

I confess at this point I don't remember if they're in this document. I'm going to take a moment here to refer to the document. Karen's on the call, if Karen wants to speak up and remind us as to whether or not she considered the Subsequent Procedures PDP recommendations, if she cataloged them at all.

KAREN SCARFONE: I believe those came out after I had submitted my report, or the current draft of my report. They're not included in the draft that you all have.

JAMES GALVIN: Actually, it is, Karen. I just found it. Section 3.7.1. It is there.

KAREN SCARFONE: Okay. Never mind.

JAMES GALVIN: Yep. You did actually pull that out and put it in there. They are definitely there. They're just not in our discussion notes, Anne. I'm going to just make a note in the discussion notes for completeness so that we also have it here. But it's definitely in the draft document so that's a good thing. Okay.

We will plan on meeting next week. Hopefully, Patrik and I will have had a chance to propose something for this discussion group here to consider in terms of capturing these two points: the need for resolver data and the idea that there are some questions that need to be answered related to that. We'll have to consider what the best way is to approach that and get that into this work product moving forward.

I know that from Karen's point of view she has done the work that she can do. If anyone happens to come up with an additional reference, there is always room for adding more items in here if they're not there. Please do send notes to the mailing list and we'll figure out what we can do with that.

But this document seems to be pretty close to being ready for review by the broader ICANN community. A public comment period, making some kind of publication there so that we can get an even broader review of the document and we're likely to get some interesting input from the community at large, too.

But I just want to prepare folks for the fact that, at some point here, we're going to do a cut-off, declare the document done, and move that along. The only thing that's missing is a bit more information on the recommendation for the future. We'll have to evaluate what the right way is to handle that and how we're going to put that together. And so, we'll seek to have more information about that next week, if not before, on the mailing list.

With that, going back to our agenda, Kim. I think it's important just to remind people that we are meeting on the Friday before the ICANN meeting. We do have the room for all day. I will use as much time as we need, or can. We will start at 9:00 AM. It will be remote participation available if it turns out you're not making it to Cancún. Just keep that in mind.

Those are the hourly spots that we have so that there's a break time in there and there's a lunch period. Our agenda is taking shape, I think, at the moment. I have already taken note that us thinking about specific analysis questions that we want to ask of data to understand the characteristics of name collisions is a good topic to have some long brainstorming sessions about. That might be a useful thing for us to do/for us to have.

In addition, there are also preparations for ... We'll have to see where we are in the public comment period and what happens there. There might be some issues to discuss there and go through that, so watch for agenda development as we get a little closer to the meeting time.

In the meantime, we will meet next week, same time, same Zoom room. Each person has their own room. Attendees have a generic one that they get to listen to. With that, any other business from anyone? I'm not seeing any hands. One last call. Then I give you back just less than ten minutes of your day. Thanks, everyone. See you next week.

KIM CARLSON: Thank you, all.

[END OF TRANSCRIPTION]