
ICANN74 | Policy Forum – NCAP Discussion Group [C]
Tuesday, June 14, 2022 – 16:30 to 17:30 AMS

JIM GALVIN:

Okay. That was nice and loud, just in case you were nodding off here. Not anymore. All right. Welcome, folks. This is our discussion group meeting. We're going to keep this relatively open and really just informal. We do have quite a few people on the Zoom. I've got a list here I can see. I'm watching carefully for hands. So if anybody wants to speak, please do put your hand up. And we'll make sure that you get in here. Otherwise, I'm going to just watch around the room here, let people wave, and/or just jump in and talk as we go along.

All right. So Matt and I both are sitting up front here. And let's get this meeting started. We decided to do something a little different than we had originally planned for this meeting for the agenda. Previously, we'd been talking about starting a tabletop exercise. And so this is an opportunity just to remind people that we wanted to dig in, now that we feel like we have at least something that's relatively stable.

And as the writing team starts to really dig in and do its part of the job of trying to write all of this out so that it makes sense, we want to jump to trying to work the process. So we thought we would take a few example TLDs. We would pick some real strings

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.

and make up a little bit around—we'd do some real data, but we'd make up some educated-based questions about them.

And we would walk through this process a little bit with the objective being to create a sample report that a technical review team might produce so that we could also include that in our final work product. So we would be setting something up for the board and for the community to look at what the TRT is going to produce.

So I did jump in here and say that we should be relatively informal and keep going. But my co-chair over here is beating me up. And I'm looking at stuff in the chat room that says, "I really just can't do that." Among other things, I didn't say my name. I am Jim Galvin, one of the co-chairs and Matt Thomas, the one over here, is reminding me to be a little better about formality.

We should at least give Jennifer a chance to do a project status update for us before we jump in and get to our work. And actually, what we should start with is asking if any discussion group members have any updates to their statement of interest that they would like to announce and present to the working group. So anybody? I'm not seeing any hands or eye-rolling. Let me turn it over to Jennifer to see if she has any project updates for us.

[JENNIFER]: Thank you. I don't have any project updates. Thank you for giving me the opportunity. I think it's best use of time to dig in. And as we had said several times in the past couple of weeks that we'll make some updates to the project found in the schedule based on how things go today and in the coming weeks. So that's all for me. Thank you though.

JIM GALVIN: Okay. Thank you for that. Okay. So we're not going to do the tabletop exercise just yet. We're planning. We never actually asked ourselves if we're going to meet next week or not.

MATTHEW THOMAS: Yes. We need to decide that. I'm planning to run it.

JIM GALVIN: Oh, ICANN's—even on Wednesday? Oh, okay. All right. So no meeting next week. All right. So we have time to figure out how to set up a tabletop exercise. Okay. So that's the goal actually for the next time that we have a meeting is to set up and start a process of a tabletop exercise. That's what we really want to do with the discussion group and work through things.

So instead, what we wanted to do is taking note of all the discussion going on the mailing list, which has actually been really very good and very much appreciate all of the interaction

on the mailing list. A lot of details and a lot of questions coming out and observations about things. So what Matt and I have done is we have pulled out of that discussion a few key takeaways. And hopefully, they are a bit of a review about where we are. But we did want to pull out of some of the discussion, some important things that we want to make sure that we are all aligned on.

These things are our sense of where we think we are. The discussion on the mailing list would suggest that maybe we're not 100% aligned. So this is the opportunity to test all of that and see if that's true and then address any concerns that come up. So we had a couple of key points that we wanted to make and walk through a little bit and just make sure that there aren't any major questions that are outstanding. If there are questions and concerns, that's fine. But hopefully, they're in the details and not a major problem.

So let's move forward one slide and we'll jump in here. So for those who were in the update meeting in the last session, this is one of the slides taken right out of that meeting about what problem are we trying to solve. We have actually said this in words a couple of times in our discussion group meeting. A version of this slide was presented to our discussion group a few weeks ago. But Matt and I had updated it and tried to get a little crisper about all of this and quite articulate in the last meeting.

So this is the same slide from the last meeting for those who were there during the community update.

And I just want to emphasize that it's easy to say that the problem we're solving is simply responding to the board resolutions. But more broadly, reminding ourselves that we're really looking for a methodology for evaluating and reducing the risk of the delegation of a TLD string. We know that we're not going to be able to solve that problem in a perfect sense. It doesn't seem, as best we can tell, that there's any kind of objective solution.

Even the 2012 round and the controlled interruption was done there, it became a mechanism for making name collisions visible and highlighting them so that there was an attempt to be able to assess what was going on and assess the impact and the harm, if you will, that could happen if that name was delegated so the board could make a decision about what it wanted to do.

The unfortunate thing is they didn't really have a mechanism for saying no. And they didn't really want to invent something at the time to just get past that issue. So in particular, we have three strings, corp, home, and mail, which are just sitting in a deferred state at the moment. And they're just kind of sitting out there. No final decision has been made by the board on those strings. So for the last 10 years, they're just hanging out.

And part of it is waiting for this framework to come around and this methodology so that we could have something which was a bit more of a process. Predictable might be going a step too far. But at least you have a process and you know how it works. And you'll have some access to data as you go along so everybody can see and it can be transparent about what's happening.

The other important part of that is really the way that we're thinking about this. And we've been saying this the last couple of weeks that we've been meeting, right? Our goal here is to identify the high-risk strings, trying to find some way to identify high-risk strings. And the default would be that strings are going to be delegated unless you can peel it off and decide that it's a high-risk string for whatever criteria or definition that we can create here.

We obviously have 10 years of experience to suggest that yes, we know that bad things can happen. But we've accepted that bad things can happen. And we used the controlled interruption as defined in 2012 as a way to identify some of those bad things. And as long as nothing crossed some line that was not well-defined at the time but they were able to identify corp, home, and mail as falling into that category, then strings would get delegated.

And we're not trying to change that model. We don't feel that we have a means of doing that. So we'll just continue to move on.

We're accepting the fact that name collisions are here. They're always going to be here. It's a risk management problem. But even as a risk management problem, it does beg the question of whether or not we can objectively identify high-risk strings. And as a category of that, can we objectively identify in advance even do-not-apply strings?

The SubPro recommendations are looking for a do-not-apply list. And it's not clear that we can create a do-not-apply list. But it does seem that we can create a methodology that allows us to peel off strings that are high risk and potentially problematic and probably should not be delegated. And that's kind of what we're looking for here. So let's move onto the next slide, please. Oh, I got a hand up there. Let's hold off and Jeff, go ahead, please.

JEFF SCHMIDT:

Hi, folks. Good afternoon. Two questions. First of all, I saw this slide earlier. I'm not sure what no other string would be blocked means. I know in the previous session you tried to explain it. But maybe it's early and my head just isn't getting around that. So if you could explain what that means a little.

And then secondly, the bolded there, that collision analysis is a risk management problem, I completely agree. And I think one of the issues that causes some of the circular arguments in this team is we have not been [crisp] on defining a threshold of

acceptable risk. Risk management is all about driving risk, managing risk down to some acceptable level and then managing or mitigating or refusing the residual risk.

We use terms like ... And I noticed in the earlier conversation there was this reference to "severe," again, from the Casey report, which I just energetically disagree with. But I understand where some may perceive those as "severe."

And so I think until we come to alignment on what is an acceptable level of risk, whether we are really seeking zero risk, whether we are seeking no reports to ICANN, whether we are seeking no glitches as a result of some delegation of some TLD, or whether there is some range that is acceptable, those ranges get really difficult, right, as we all know. A thousand companies have a problem. Is 1,001 okay? Is 10 okay, etc., or there's other ways to express that. In 2012, we used risk-to-human-life as a threshold.

But anyway, I would encourage this group to consider very deliberately what is acceptable in the wonderful and messy Internet, in the innovative Internet, what level we would declare success.

One way to do that is to reflect on the 2012 round and say, "Was there any unacceptable risk that surfaced or any unacceptable behaviors or glitches that surfaced based on the 2012 processes

and criteria? If so, what are those and what can we do in the future to manage those down?" Thank you.

JIM GALVIN:

Thank you, Jeff. Excellent questions. Let me try to give a very couple of quick responses and get some other voices in. I see we now have a queue of hands. On the question of what does it mean for the default to be that the strings are to be delegated, the idea here is if you apply for a string and if you pass all other criteria which are not within our scope and that we're not addressing, then your string will get delegated.

Frankly, that is essentially what happened from the 2012 round. They delegated everything. And unless something bad happened, you continued on to have the string. So they granted you the string, and that that was that. So the idea here is we are looking for reasons not to grant the string. Otherwise, it will be granted as long as you pass all other criteria. So that's what that second sub-bullet means. And we'll try to find a better way of saying all of that when we write the document.

On your comment about risk, one thing I want to call out for us to keep in mind is it's not our job to decide what level of risk is good or bad. It's our job to look at what's going on and to explain from a technical point of view the issues that manifest if this string delegates or not. That's the role of the technical review team.

And one of the things that I would hope that we will get out of doing a tabletop exercise where we ourselves pretend to be a TRT and we work through that process, we want to identify what are the things that can happen. And then, the board gets to evaluate those and decide whether those things that can happen are too risky for it. And then, it can decide whether to grant or not to grant based on that. It's not really our choice. I'll give you a two fingers if you want to jump in. And otherwise, let's go down the queue here.

JEFF SCHMIDT:

Thanks for that. So first of all, on the no other string, I think what you're saying is the default is to allow, right? But for compelling evidence to deny, the default is to allow, and that makes sense. I have just a practical concern about this whole idea that we can wheel a pile of data to the board, which is basically going to say, "Bad things might happen. We're not going to take a position on whether this string is 'safe' to delegate or not and rock and roll, go to it, board. Make your call."

I think that's a mistake. I think that's going to lead to ... I don't know. I don't know what kind of behavior that's going to lead to. But I think we, as the technical experts here, are going to need to develop a process that leads to a recommendation. Now, the board can always override the recommendation.

But I think it is incumbent on us to wind up telling the board, "We think this is okay," or develop a process. Not we, but develop a process that says, "We think this is okay," or, "We think this is not okay and here's why." Again, they can always overrule us. But I don't think it ends in anything other than tears if we just say, "Here's the data. Go to it." And I'll stop. Thank you.

JIM GALVIN:

Okay. Thank you. I am hopeful that the tabletop exercise will get us to a more practical place in being able to respond to that question. So I take your point. Thomas, go ahead, please.

THOMAS BARRETT:

This may be on a later slide. So let me know if it is. But I wonder if we're being too narrow in defining who the audience is. Is it just the ICANN board we're trying to provide a methodology for, or is it applicants, or is it the ICANN community? My impression was applicants were going to be provided some information as well even prior to this issue reaching the board. And I wonder if that's defined better elsewhere on a later slide or if this is too narrow.

JIM GALVIN:

No, that's an excellent question, and you're right. We have talked about that before. And for the purposes of even the community update, we consolidated that into the bullet that

says, "That will be happening." It's inside the bullet of the applicant creating a mitigation or a mediation plan. And we just did not expand that detail into if they're going to do that, they have to collect data.

So I think that we'll shift that to being part of the study three investigation, not part of this. That's the current working model. But you certainly can continue to challenge that for now as we fill out what we're doing. Okay. Thanks. Rod?

ROD RASMUSSEN:

Yeah. So I just want to reiterate, and I guess it's two fingers. Our job is not to determine what the board's appetite for risk is, right? So that's out of scope. We could probably offer some suggestions or helpful hints as to—and that's part of our job here—if you do this, this will probably happen. You need to decide whether or not this is a risk you want to take. So it's not just throwing a pile of data at them. There should be hopefully some guidance around what are the kinds of things to expect from this type of thing that we see and that you see in the data can turn into this.

But that would be, I would posit, for whoever's doing the analysis to take on and provide to the board with an assessment of, based on things that we can provide some guidance on from here, as to, "These are the things you may be exposed to," you being global you. Who knows what. It's going to be various

people could be exposed to various things. But as SSAC chair, I can definitely say the SSAC is not going to tell the board what level of risk it's going to take on or not. And that would not go very far. Thanks.

JIM GALVIN:

Thank you, Rod. And I have to say I agree 100% with what you said. That really would be the TRT. And you're right. I take Jeff's point about it's hard to imagine that we wouldn't give a recommendation. But a lot depends on how you write the work product, right? Yes, the TRT is going to have significant influence on the decision that the board makes.

And how they write and describe the risks is going to matter. And that's going to take shape as we go along here. And I guess we'll examine that as we do our tabletop exercise and get into it better. Thank you. So Warren, hand up next.

WARREN KUMARI:

Thank you. So I think I'm sort of agreeing with Jeff and Rod, which is a somewhat odd position because they have sort of said opposite things. I think that we've been doing this for an awfully long time now, the NCAP work. And if I try and summarize it, we have the people will look at some stuff and decide to apply or not. And then, there'll be some sort of testing. And then, there'll be a decision made. But it feels like there's a huge chunk

between that and how one actually gets to any sort of useful risk assessment.

We seem to be doing a fair bit of “and the TRT will decide” —sort of punting the problem over to the TRT but without there being much sort of protein or meat on the bone of what people actually are supposed to use for the risk assessment. The closest thing that we have to any actual guidance or data largely seems to be people could maybe look at the magnitude stuff, which I fundamentally think doesn't really do a lot as a flawed data source.

So people could maybe look at that. And then, there'll be some sort of testing, passive, active, whatever. But there's a huge, or it feels to me, like there's a huge amount of handwaving still around. There'll be some sort of testing and analysis done. And there'll be some data collected. But we don't necessarily seem to be making much traction on what actually any of the details are.

And “we leave that up to the TRT” feels like a fair bit of we haven't done a huge amount of the actual work. The first and second bullets are the ICANN board needs a methodology for blah, blah, blah. And what we're saying is we should propose a methodology.

But we don't actually seem to have already proposed the methodology yet. I feel like I might be somewhat of a downer

here. But we have the document that has a pretty flowchart. But there's a lot of “Here be dragons. We will figure that out later” parts. And I think I might have annoyed Rod.

JIM GALVIN:

No. Thank you, Warren. I think I'll just offer two things. I take your point, and that's a fair comment. Ultimately, as we finish getting a lot more down in black and white, this will hopefully feel better. And as you're part of the writing team, I'm looking to you to make sure that we stay on track and on that point and fill out the details in a way that satisfies your question.

It's a good question to keep in mind and always have there. And I appreciate you've brought it up before. And you should continue to bring it up and until you feel satisfied. And we're going to work on that. My answer is that I am hopeful that as the writing team gets a more complete document for people to look at, this will begin to look better. And we'll be able to more carefully identify the questions that are not being addressed instead of the broad-sweeping question of just not getting traction yet. Maybe we can identify the real specific gaps. I am hopeful that we'll get there. So thank you for that. Okay. Next hand is Anne.

ANNE AIKMAN-SCALESE: I'm sympathetic to the notion that the criteria aren't hard and fast yet. I think my understanding is that it's the CDMs that end up forming those criteria. It strikes me that as to the board decision from the 2012 round, of course, the board relied on a device from a neutral third party. And they had I think .corp and .home on the list. And then, from the JAS report, they added a new name collision occurrence management framework. They had decided they would add .mail at that time.

And I'm interested that to make it possible that our result for this process and how the technical review team operates be as independent and neutral and objective as possible because I'm hoping that when reports come out of the technical review team to the board that they're consistent enough in terms of the determination made by the board that the board itself doesn't run too much risk of requests for reconsideration.

So the criteria I think have to be as objective as possible. And compared to Jeff Schmidt's indicated well, we did all this when I think before we made our report of whatever back in 2012, I guess I'm interested in making sure that the board has the same feeling of confidence in the decisions that it makes about risk—that the information coming there ... And I know I'm being repetitive, and I apologize.

But the board's going to end up delegating that to its technical committee, right? And the technical committee is going to need

to make decisions that are consistent. And so I'm very supportive of the notion that the criteria that we list in the CDMs be as objective as possible.

JIM GALVIN:

So thank you for that. Let's hope that we get there. You were tapping your two fingers? Okay. Well, go ahead, Rod. We did have a couple other slides I want to get to. Go ahead.

ROD RASMUSSEN:

All right. I just want to maybe belabor this a little bit but so we've got CDMs, which we're working on, right? And that's subjective criteria of some sort. And they may or may not have enough protein. We don't know yet, right? We're driving towards that. Those then create, in my mind at least, the process goes as it is. TRT will objectively look at this through the process, get data in, and look at that and see, "Okay, what is this telling us?" That's where we can provide at least some guidance here as to what are the things to look for, etc.

But then, there's the practical part of, "Okay, why is this happening?" And we already talked about this from a mitigation perspective. We can take a look at things. "Oh, dot [inaudible]. Well, I guess we know what that's all about. And let's confirm that." And then, there may or may not be a recommendation for that. It will be different than a corp or a foo or whatever.

And that's where it's really hard to say in set, hard concrete, "Here is the decision or the recommendation to make," based on actually looking and trying to as best as possible understand what's causing a CDM, a set of CDMs, to hit a, "Hey, we should look at this parameter," right? And I don't know that we're going to recommend to anything saying, "There's a threshold number." There's absolutely no way you can do this is. Look at it and try and understand what's causing that before you make any kind of recommendation.

So it just gets really hard without getting into practical examples of each string, looking at them uniquely, as to how you would process it. But if we can provide guidances to going through a process and thinking about mitigation, all those kinds of things, then we can at least provide guidance to what a TRT could do with it as objectively as possible. And if we strive for that, that's about as good as I think we can do.

JIM GALVIN: Thank you. Thank you for that. Excellent points. I tend to agree with you. Okay. Let's move on. Oh, Warren's got his hand up too.

WARREN KUMARI: So just to respond to Rod, yeah, I fully agree. I don't think that we should say seven is an acceptable number, regardless of what scale we are using seven on. But we don't yet have

anything like the CDMs. We keep saying, "We should have those." And then people can make an evaluation based upon that, once there's some stuff. But it feels like this is kind of where we were in 2012 with some, "We should do some analysis of some sort of stuff and come up with some way to evaluate." And we haven't actually ...

JIM GALVIN:

So that's a nice segue into the next couple of slides. Okay? Let's go to the next slide. And we'll very quickly look at this. You really have seen this slide before in our updates to the group here. I'll just highlight the first bullet up there. Our goal here is about ensuring that name collisions can be assessed and making them visible. And that's the big part of what PCA and ACA are all about.

And, of course, we do have a goal of wanting to make sure there's an opportunity for a mitigation or mediation plan. We've kind of collapsed that down to the end of step four, which is active collision assessment. And we're going to really push most of that work into a study three. But that's the spot where that would be. We're focusing on the first bullet there of ensuring name collisions can be assessed.

We already know about the two operating roles that are needed. I'm not going to go back over those here, the TRT and the neutral service provider. Let's go to the next slide here. And this

is coming right off the mailing list and discussion. And I'm going to turn this over to Matt to walk through.

MATTHEW THOMAS:

Thanks, Jim. So this came off the mailing list from Jeff Schmidt. I hope I accurately represented your e-mail here, Jeff. I copied and pasted the high notes. So I think it should be fine. But I think the higher order bit here that this e-mail thread was coming back to was defining the management program in terms of what we're talking about today within the discussion group versus what went on in the 2012 round.

And to your point was what can we improve and maybe what things went wrong in 2012 or what we could do differently. And so I think this is an important opportunity for us to kind of level set and think about what some of those things were that we've discussed about here in the discussion group.

I know we've talked about it in the past that in the 2012 round that there was a whole bunch more contextual items around how NCAP was part of the round of TLDs. And it wasn't something that had a greenfield opportunity like we have here within the discussion group to kind of look at it more holistically.

There were also concerns around particular portions of when controlled interruption was being deployed because it was done essentially when the applicant was granted the TLD. So I think

there is a few key components here that we should examine as a discussion group in terms of why we think those things should possibly be changed into this new framework and workflow that we're discussing and why they're appropriate.

But I think at a whole, Jeff, I think we're largely kind of aligned with what the process was in 2012, right, with our current workflow thinking. There is a need to identify these high-level strings or high-risk strings, which you've termed black swan strings. And then, having a mechanism for notification and remediation in there and then, ultimately, having some kind of emergency response.

So I think there's a large, broad overlap between what was in the 2012 with some minor additional tweaks and provisions in which the new system that we're proposing addresses some of the concerns that we've observed from the previous round, as well as address some of the data concerns that has happened with the DNS data available and its fidelity in the root server data and the general DNS ecosystem evolution. I just wanted to kind of throw that out there. Any thoughts from the discussion group or hands on that?

JIM GALVIN:

No.

MATTHEW THOMAS: Okay.

JIM GALVIN: I'll just add, in particular, calling out Jeff directly, part of what we were trying to get out of the discussion and a key takeaway was really to make sure that we all are aligned and understand that we're not replacing what was done in 2012. We have simply discovered a couple of key areas in which we can evolve it. The concept is the same. We're going to do an interruption.

We're just breaking it down into multiple steps before we get to the final interruption that really does change the behavior for clients. And we're allowing for the possibility of adapting the CDMs that were done in 2012 to some new protocol. So we want to have some discussion about that as we continue through this. We do think it's important to add some additional information.

And then, of course, tweaking the notification idea. The 127.53.53 idea was interesting at the time. But the root cause analysis document has suggested that it was not as effective as we would have liked it to be. So given that, let's tweak this a little bit. And let's move forward with something a little different and something which is a little more directly responsive to the client rather than just being responsive in a log file to a system administrator.

So broadening the scope of visibility of the notification would seem like a good step to evolve to. And that's really all we're after here. Okay. I'm not seeing any hands, comments, questions. We are a quarter after the hour. We have 15 more minutes in this session. But really just looking for some affirmation that all of this sounds like what you were expecting. It sounds like what you understood. Oh, Anne, you've got your hand up. Let me just quiet down and jump to you.

ANNE AIKMAN-SCALESE: Thanks. Just a question, and again, in comparison to what occurred in 2012. So at the time, there was this whole thing of alternate path to delegation. And then, ultimately, in the framework that was adopted, there was controlled interruption. And it occurred 90 days. And this occurred after the contract award and after delegation.

In connection with our active collision assessment section, I guess I'm assuming that the active collision assessment happens before the actual authorization for contract award or whatever. And I don't know if that's a correct assumption.

JIM GALVIN: That is the current working model. And just to be clear, I think that even originally for me when this group first started, and Matt came along partway through, but our just assessment of

this and our interpretation of what's going on is that really is what the board was looking for was a methodology to deal with name collisions and assessing that risk.

And I agree. I don't think it was ever stated explicitly. But it just felt like the presumption was that all of this should happen before you grant the TLD. They made the choice that they made for 2012, which is fine. But I think from a technical point of view, we're getting to a place where you really should do this first. But ultimately, I guess the board gets to decide which is harder or easier. Is it harder to take it away once you've granted it, or is it harder to do this assessment before you grant it?

I suppose from an implementation point of view, they could shift it and make that different. But our going-in position here is to do this in front. And then, we'll just see what they do with that. Yeah, Rod, you've got your hand up. Go ahead, please.

ROD RASMUSSEN:

Yeah. It's out of our scope. It really is. Yeah, it's logical. We could opine on it but it's not really anything we were asked about. And frankly, that's a decision for us as an implementation that has ... If somebody could come up with a really clever reason as to why technically or from an impact level, whether or not we do it, then we could argue that way from that perspective. Great, we could do that but off the top of my head, I can't really think of one.

ANNE AIKMAN-SCALESE: Yeah, I see the point that it's up to the board. But I guess what I'm trying to understand is apparently, if we are trying to determine strings that are so high risk that maybe they shouldn't even be delegated, and if that process is continuing through the active collision assessment phase of what we're doing, which I think it is and it's quite consistent, by the way, with the SubPro recommendations, then you kind of have to make a recommendation with respect to when that process occurs because if you're going to determine during active collision assessment that the string is too high-risk to be awarded and delegated to move on, then what do they say? The horse is out of the barn if you've failed to make that assessment prior to the award.

ROD RASMUSSEN: There's a difference between award and delegation that way. Clearly, from a technical perspective, delegation to a third party is ... That's within our agreement to opine on. The award, that's a business risk more than a technical risk of any sort.

ANNE AIKMAN-SCALESE: Just real quickly, and I think it's a really good point, but I'm assuming this whole thing happens before the rest of the evaluation of the string. And the reason I'm bringing this up is

because, again, of that SubPro implementation guidance about 29.5 that says that we should construct a system where the applicant—that they don't need to proceed through evaluation if they've been tagged as a black swan. By the way, I think black swans are beautiful. But if they've been tagged as that, as a high-risk string, that they may want to withdraw and not proceed through the rest of the evaluation.

JIM GALVIN:

Okay. We have some other hands up here. I just was thinking about something about the gaming thing. But let's run out the queue first here. We've got 12 minutes left. Go ahead, Steve Sheng.

STEVE SHENG:

Yeah, thank you, Jim. So looking at the risk management, look at the slides here. In the 2012 round, the identification of potential black swan strings and also that is done purely using the DITL data. And what's being proposed here is a temporary delegation. So I think in the report, the case needs to be made clearer the benefits of that. And I think a lot of the data points are already there. But I think it's useful to make that case clearer and strong.

The other thing, regarding number one, the temporary delegation would require some actions by IANA, obviously. And

there may be hurdles that needs to be checked ahead of time, right, because for ICANN, for anything put into that root zone, that's a big step for ICANN.

So I think some conversations needs to be had. If IANA were to do this, what kind of process it needs to go through because I think it would be unwise for the working group to come in the end when it delivers its report and for IANA to say, "Well, you haven't considered this." So I think that's worth checking.

The second thing really, using live data through a temporary delegation, and as Warren showed so aptly, is it's so easy to be gamed. I think there needs to be some dedicated discussion in the report on that because that's a big issue, especially when live data is used. Thanks.

JIM GALVIN:

Thank you. All excellent points. We'll have to make sure to capture those so that we can speak to those issues. Let me go to Warren.

WARREN KUMARI:

So thank you. Steve actually set it up really nicely for me. I think we need to discuss the gaming problem a lot more and sort of the discussion of using the magnitude list. If I started with no resources for less than \$200, \$200 per month paid in Bitcoin, I

could easily get to the very top of the magnitude list. And that's just renting a server in Russia and sweeping queries.

There's a bunch of risks in that. One would be an applicant gaming their own string and then showing, "Look. Somebody's gaming me. You should please put me on the alternate path because obviously, the data is now tainted." There's the if a string looks like a black swan and the applicant is supposed to withdraw, there is a bunch of incentive there for their competitors to try and make them look like a black swan.

But apart from all of that, I think there is significant risk to the root system itself if people start trying to use this list for something because any gaming is probably going to be at a fairly high volume. And it's going to be a bunch of queries that can't be answered and sort of get cached [really well]. So it's not just a how badly does this end for the applicant. It's also what are the collateral damages that could happen from people trying to game.

And again, that's really, really easy to do. You rent a server that allows spoofing for \$200 in Bitcoin or many people already have infrastructure. If you're going to apply for something, you probably at least have some computers. And you just start making up queries. And very much by definition it's going to be hard to tell if they are gamed or not because if we could tell that,

we could have solved most of the name collision problem already.

So, again, I think we need to be really careful about talking about using the fact that a string shows up at the root as being evidence that it's definitely a bad thing or is okay or isn't.

And then, there's also the thing of there's a huge amount of stuff which the root is not a useful vantage point for because of aggressive NSEC, and caching, and local root, and things where there's an authoritative thing, an authoritative server so the names don't leak out. And so I think there are a lot of sort of technical issues that we need to be keeping in mind about what is a useful data source and what isn't and what data is actually visible or what isn't. Soap box rant over.

JIM GALVIN:

I agree on all those points. And I also agree on we certainly should say what we know about gaming—what's hard and what is potentially available that TRT should look for, and then the kinds of things that will be hard to detect because all of this we'll learn over time, right? Maybe over time, people will find different ways to do things. They have to evolve. The TRT has to evolve as well as the bad guys have to evolve.

Gaming will always exist. So we certainly should document what we can about all we know and how to deal with it today and

allow for the possibility that TRT will have to learn as it goes and do a better job as it goes along. So we'll certainly have the discussion about gaming in our document. I just don't know that we will solve that problem, as you said. It will just be something that's visible. Go ahead, Suzanne.

SUZANNE WOOLF:

Sure. Thanks, Jim. I just want to call out what Warren just went through. When we're talking about gaming, it's one thing to say we're talking about some obscure and hypothetical attack on methodology. I want to emphasize what Warren is pointing out here is an extremely trivial way. So I think there's a question of degree there because if anybody remembers the, they called it digital archery that was supposed to be a mechanism for determining a contentious set.

I think it was a long time ago. I'm fuzzy but what I remember very distinctly was that when people that were working on those issues realized how easy it would be to game, people sort of looked at each other and said, "If we go with this, not only are we going to get gamed, but we're going to look stupid." So I think we can't rule out all possible hypothetical sources of attacks. You're right, nothing's perfect.

But at the same time, if even sitting in this room we can find an easy way to hack on a proposed test or data source, I think that's worth taking into account. And I know that makes life harder for

everybody. But frankly, we don't want to look silly. And we don't want the community to look silly. And I think there's a question of degree.

JIM GALVIN:

I agree. My comment about the gaming in this one particular scenario. So Warren was talking about, "Gee, I can get myself on the top of that list." And so my first reaction to that is well, just getting on the top of the list should be fairly straight-forward for the TRT to see because you can look at L-Root historical data.

And if it wasn't there before and it's there now, it's like, "Oh, okay. This is interesting. It wasn't there a month ago. And now, it's there now." Or you go back through some amount of historical data. And all of that has to be sorted out. But that's really the idea.

So getting on the list is trivial. Now, the question is what value is being on the list? And what are you trying to achieve, right? And that's where you get into this degrees question. And I agree with you there. And as we get towards three minutes before the end here, I guess I'll stop talking about gaming. But Warren looks like he wants to add something. I'll let him jump in here.

WARREN KUMARI:

So number 20 on the list is .dummy. You have no way of knowing that I haven't been putting myself as .dummy for the last while

because it's been clear since 2012 there's going to be a new round. So looking back a little bit doesn't necessarily show stuff. But also, if we just did the, we'll ignore gaming by looking back, then it ignores all of the things like the [net playback], which is towards the top of the list at the moment and .console, which showed up and was a major issue at A-Root for a while. And it just appeared, and it would have been as significant problem if it had been delegated.

And there's openstack.local. It didn't exist, and then suddenly did. And that was a large number of people started deploying OpenStack software. So either you assume that nobody has been gaming until a specific time and anything after that you're going to assume is gamed, or you need some better way of figuring out what's gamed and what isn't.

And I think it's really easy to generate data that is basically indistinguishable from gamed because all that a server sees is a DNS query. So I can easily spoof those. I can easily tell what looks like a real DNS server. And so I think that the real way to detect gaming by looking backwards doesn't really work in order to detect it because it looks synthetic. It also doesn't work. So anyway.

JIM GALVIN:

So I agree. I am not at all trying to suggest that we know all of the objective criteria for deciding that something is gamed or

not. And I didn't mean to suggest that looking back is going to answer all the questions. Those are issues to be considered. As in all of this, looking back is simply an indicator. Whatever the number is in terms of how far back you go, 1 year, 5 years, or 10 years, you're right. You don't know how long somebody's been gaming or planning for this. But at least it's a step in the direction of things that you look at.

And then you go look at other things. And you're right. Eventually, you actually have to look into where those queries are coming from and why. And there's a limit to how far you can take that. As we discovered in the root cause analysis that we were trying to do with Casey, he discovered quite quickly that there was a limit to how much root cause he could actually determine. He could only go so far.

And, in fact, he ended up not being able to go very far at all. Of course, in his case, he was looking at really old data to examine. Maybe things will be a little different if it's more current. But then, again, you still have to believe there's a limit. So it's not perfect. But people are going to have to do investigative work when you get down to it. And you're going to have to make a judgment call about whether or not how suspicious things look or don't look.

And that's what the TRT would have to do to make that judgment call. We can try and lay out the framework for all of

this in guidance for dealing with that. And Anne, I'm going to give you the last word here. And then, we will close ourselves out.

ANNE AIKMAN-SCALESE: Thank you. Just something that perhaps is an implementation detail but you know there exist in name collision occurrence management framework. What we're coming up with essentially would I think replace that. And it would have to be drafted accordingly.

And I'm assuming that what we'd be saying in a final report is that that drafting exercise is up to implementation team or the staff or something like that. But a framework exists and what we're developing is actually a new framework.

JIM GALVIN: Yes. And as you would well know, even in our case, we can make implementation notes in whatever we want to say if we want to influence a direction, as long as we provide some justification for that. So that certainly is a door that's open to us if we want to go there just to be helpful. Okay. With that, let me say thank you to everyone who joined us, both remotely and here. And we will not have a meeting next Wednesday. And watch for more information on the mailing list as we move forward. So thanks, everyone. We're adjourned.

[END OF TRANSCRIPTION]

CLOSED SESSION