DANIELLE RUTHERFORD:	Hello all. Welcome to the RZERC Teleconference held on Tuesday, the 18th of May 2021 at 19:00 UTC. Tim, would you like me to start the roll call?
TIM APRIL:	Yes, please.
DANIELLE RUTHERFORD:	Carlos Martinez, ASO? No, he's not on the call yet. Peter Koch, ccNSO?
PETER KOCH:	Yes, here I am.
DANIELLE RUTHERFORD:	Kaveh Ranjbar, ICANN Board?
KAVEH RANJBAR:	Yes, present.
DANIELLE RUTHERFORD:	Kim Davies, PTI?
KIM DAVIES:	Present.

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.

DANIELLE RUTHERFORD: Tim April, IETF?

TIM APRIL: Present.

DANIELLE RUTHERFORD: Howard Eland, GNSO Registry Stakeholder Group?

HOWARD ELAND: Yes, ma'am.

DANIELLE RUTHERFORD: Daniel Migault, RSSAC? I note he's not on the call yet. Duane Wessels, Verisign as the Root Zone Maintainer?

DUANE WESSELS: Yes, Duane is here.

DANIELLE RUTHERFORD: And Geoff Huston representing the SSAC.

GEOFF HUSTON: Yep. I'm here.

DANIELLE RUTHERFORD: All right. Thank you. Tim, I will pass it over to you.

TIM APRIL: Thanks, Danielle. I just saw regrets from Carlos a minute ago in the thread. The first thing is to review the agenda. Any comments on the agenda distributed last week? Okay.

Then into the minutes from last week—you can pull that up, Danielle. Everyone okay with going forward with these minutes or any objections? Good. Thumbs up from Peter and Duane. Okay. Thank you.

The next topic was discussing RZERC002 and 003. Paul has been kind enough to join us for the discussion about this. At the last meeting, we were discussing the Board Statement of Understanding related to RZERC002/003 and we were going back and forth on a bunch of them and figured it might be easier just to have someone who was in the other side of this process here to talk through what would be useful in replying to this.

I think the overall, from my understanding of what we discussed last time, was that in most of the cases we were trying to amplify RSSAC028 and basically trying to get this work prioritized. They're trying to defer most of the decisions to RSSAC for clarification. That was covering I think the first four comments on 002, if I remember correctly.

DANIELLE RUTHERFORD: I think so. That covers 1(a), 1(b), 1(c), and 2(a).

TIM APRIL: I guess the first question I would have to you, Paul, is would thus saying we would defer to RSSAC for clarification on how each of those would be judged be meaningful and helpful to you and OCTO for this assessment work?

PAUL HOFFMAN: Yes. I don't know if you're able to do that. I don't know if that's allowed. As I understand at higher level, RZERC is asking OCTO to go ahead and do Recommendation 2 of RSSAC028, which hasn't been done. RSSAC028 is out there. Recommendation 2 is sitting there and no one has done anything. My understanding overall is RZERC is saying, "Please enact that." Is that correct?

TIM APRIL: Yes. That's my understanding. Peter has his hand up.

PETER KOCH: Yeah. Thanks, Tim. Hi, Paul. Paul, you said on a higher level RZERC is asking OCTO. I think on a higher level we are asking the Board, and then the Board decides who to task with it.

PAUL HOFFMAN: Yeah.

PETER KOCH: Just because we were on the higher level.

PAUL HOFFMAN: Yeah, fair point. By the way, going down with the next level, once the Board says, "Yes, thank you for saying that," they're going to turn over and look at OCTO, and OCTO is going to turn over and look at me. I don't know to answer your question, Tim. I don't know if it would help for you to say, "We defer to RSSAC." It may be fine just to leave it alone. And as we are progressing with those—which turned out they were not easy questions. One of the reasons when we did 028 that we punted on them was it wasn't clear necessarily how to those and things like that. It may be that just as a normal part of ICANN, OCTO, me—or actually probably find a subcontractor to do it—working on this is that we would have questions about how is this going or whatever, and then we could either turn and ask RSSAC or we could turn and ask RSSAC and RZERC. I don't know if you want that kind of interaction or not on the RZERC level. We certainly would ask RSSAC because they are RSSAC recommendations, but it seemed like you folks cared about those. So do you want us to also ask you as we are doing them when we come up with questions or implementation questions and such like that?

PETER KOCH: I think that's appropriate, personally. But it may depend on the nature of the questions, too.

PAUL HOFFMAN: Sure.

PETER KOCH: But I think communication is definitely appropriate as the work progresses.

PAUL HOFFMAN: But my question is should we be thinking that our communication would also come to RZERC as well as RSSAC?

TIM APRIL: Kaveh?

KAVEH RANJBAR: I think it's a valid question. I think it's better if we are informed. I don't think we need to introduce that much formality. So my suggestion is OCTO to please keep us in loop. Inform us what they're asking, what you're receiving, and what actions you're planning to do. And of course, if you really have an issue, we can discuss within RZERC and then we have liaisons to RSSAC and RSO. So we can give that message formally to OCTO, if needed. But I think it's simplifying the process because there's a good chance that there's no real thing here is to keep RZERC informed, please. But it doesn't need to be stopped in the process. I think RZERC has enough ties to RSSAC to be able to …

DANIEL MIGAULT: One point I'd like to clarify, is the idea to mostly involve RSSAC? I mean, having questions to RSSAC and then keeping RZERC informed or is that the other way around?

DUANE WESSELS: I think it's the first way, Daniel.

DANIEL MIGAULT: Okay.

PAUL HOFFMAN: Yeah. That was my question. What I heard was, what Duane just said, which is that we are asking questions of RSSAC if they come up—and I strongly suspect they will—but we will keep RZERC informed as we're going through. Given that RSSAC028 had one recommendation with four bullets, I still think it would be appropriate if the end of this project was all four bullets are done. That is that we wouldn't do incremental, like we wouldn't split them up. But we could let you know, "Oh, we think we're at top of the stack on this one, sort of where we're at," and see if that's helpful.

TIM APRIL: I was going to agree with the others. I personally think that having RZERC informed but defer to RSSAC on all of the technical questions and things like that seem to make sense and there's no overlap here that I think that works pretty well. Steve?

STEVE SHENG: Thanks, Tim. Just hearing the conversation, what I would recommend is, when OCTO starts to doing the work, you can update the progress

through the ARR. That way, you only need to update once and both the RSSAC and the RZERC can be informed. There are points of time where, Paul, you feel you need some of this note in put, maybe that's to engage RSSAC first and then also to see if there's interest from RZERC. So that will be the path I'll recommend. Thanks.

PAUL HOFFMAN: Thanks, Steve. I didn't realize that ARR existed after the Board had said, "Go do something." I didn't realize that steps later were noted there. So I'll make sure that we do that.

DANIELLE RUTHERFORD: I can speak to that a little bit. When things go into Phase 4 and they're being implemented, you would receive monthly e-mails calling for updates for the ARR Team to deliver to the different Review Committees or Advisory Committees. They can handle that for you and for us.

PAUL HOFFMAN: Okay. Thanks, Danielle.

DANIEL MIGAULT: I have one question—what ARR means exactly.

PAUL HOFFMAN: I'll let Danielle answer that one.

DANIELLE RUTHERFORD: The ARR stands for the Action Request Register, sometimes called the Board Advice Register. So when any of the ACs, the GAC, ALAC, RSSAC, or SSAC do issue advice to the Board, the ICANN Org has a team that takes that in, assigns that to people within the Org for the understanding process, the evaluation process. They're in charge of bringing this item before the Board. Then once the Board decides to implement it, it would go into a phase where we get monthly updates on any advice items. So it's just a tracking process for all advice to the Board. I can put the public link where things are tracked. It's essentially a register where all advice items are tracked through open implementation and closing.

DANIEL MIGAULT: Okay. Thanks. I was just wondering the ARR, what the … The other thing I just like to raise is that I think RZERC and RSSAC mailing list do not have the same status. RZERC is public I think, while RSSAC is not. It's just a note I'm raising.

PAUL HOFFMAN: That's a fair point. Duane, can I ask you sort of a tricky question here in public?

DUANE WESSELS: Okay.

PAUL HOFFMAN: What if I post questions and such to the RSSAC Caucus mailing list as compared to the RSSAC mailing list? It's an RSSAC document but there

was caucus interaction on this. Does that sound like a reasonable way of having questions and comments and such be made in public?

DUANE WESSELS: I think that's appropriate. The RSSAC document was RSSAC Caucus's product so I think that's entirely appropriate.

PAUL HOFFMAN: Great. So, Danielle, that should mitigate that issue for you, is that I'll just do that on the RSSAC Caucus mailing list which is public.

DUANE WESSELS: I think we need to talk about the text that peers highlighted and commented on the Recommendation R1(b). I think we talked about this a little bit last time. The draft clarification has some numbers. Paul, I don't know if you're able to see this or not, but the question was, "What does RZERC consider a lot and significant?" I think I had originally written this and I just made up these numbers just as a discussion starting point. That's where they came from, Peter. I think we talked a little about this last time but I don't remember what the conclusion was.

TIM APRIL: Peter has his hand up. Actually, Steve's hand was first.

PETER KOCH: Steve, are you going?

STEVE SHENG: That's an old hand, sorry.

PETER KOCH: Thanks, Tim and Duane. Maybe I should clarify. And apologies, I did these comments in the 30 minutes before the meeting. That was a very short notice or the last moment. I wasn't sure whether we discussed it or not. I'm not so much questioning the numbers but in the context of this formal exchange. We might want to be careful so that these numbers aren't cast in stone over the course of time. Essentially, this boils down to how many kids can we leave behind in a way. That is, of course, the decision to be made very carefully. It may be close to the discussion that we saw during the KSK rollover. What is the acceptable threshold? What is the number of casualties? Maybe I should tone down my language, apologies. How many resolvers or how many people behind resolvers, actually—and I see Geoff already has his hand up so I'll try to shape that in. How many people behind resolvers are willing or is the system willing to lose or to impact? I think it's useful to have some numbers here to give some general understanding but I would not want our advice here to be read as, "Yes, RZERC said that this is the acceptable number for damage," where we don't know the extent of the damage. I hope that clarifies that snarky remark. Thank you.

GEOFF HUSTON: Look, I applaud the effort to try and put a number against what we call significant and a lot. However, I also weep at the fact that this kind of

effort is somewhat doomed. I don't think the number hits the target and I'm not sure any number could.

Fragmentation is a very, very difficult topic in terms of trying to assess the impact. For example, Google's public DNS does not accept fragmented v6 packets at all. So for one very large open DNS resolver used by 20% of the world's users, if it resorted to priming queries—and it doesn't, but if it did in a hypothetical case, the tiny percentage of "I don't like fragmentation when you're doing this over v6" affects an awful lot of users. Whereas if they're a bunch of extremely small—let's call them CPE-based systems that don't accept fragmentation that have a client count of one, then the actual impact in terms of users is entirely different to the numbers that you'd see in the response for priming queries. So in some ways while the effort to try and quantify what we talked about as a lot and significant is entirely applaudable, I think any number that you'd pick out—20%, 5% or anything else—is kind of difficult and, in some ways, completely meaningless. It doesn't correlate to users in any shape or form. Fragmentation is a much more subtle problem than that, I would have thought.

The other thing I would note is that in some ways it's not the fragmentation that concerns anyone. The priming query is not time-critical. You're allowed to take some time. You're allowed to do fallback to TCP. You're allowed to do that because no one or no application, in theory, is sitting there with a ticking clock. Once the resolver is up and warmed up, timing becomes important. But before that, I actually think it's a different problem. So while I applaud the fact of a lot and significant is being a fine thing to ask for, I really don't understand how to put a number on this and I don't agree with the

arbitrary numbers of 20% and 5%. I just don't see that they have any meaning in terms of user impact. Maybe if that's what we wanted, we should think about it again because attempting to quantify doesn't really work for me. Thank you.

TIM APRIL: Paul?

PAUL HOFFMAN: Thank you, Geoff and others, on asking this. Thank you, Duane, for picking up numbers from the air because at least now we have some numbers. But I note that RSSAC028 did not have any numbers and it didn't even start to get there, not only not in the recommendation but in the body of RSSAC028. Duane or Kaveh, actually, correct me if I'm wrong, RSSAC itself has not even wandered towards picking any numbers in either of these realms. So if RZERC wants to, sure. I'm here as staff. We will support you how we can. I would propose, though, that a way to go is that ICANN via OCTO does Recommendation 2 of 028 comes out with some numbers to answer the four bullet points so that we have some ideas. Then at that point, RZERC or RSSAC could say, "Do we think that there are important numbers? Geoff would immediately chime in saying no, which is fine. Others might chime in and say yes, anything above zero is bad. That discussion could happen organically after we've gotten the best numbers we can for what's going to happen with various authoritative name servers.

We can't simulate very well the resolvers that we care about hitting the root servers. What we know is by looking at root server traffic that a

bunch of those resolvers are not the ones we care about and it's hard to differentiate. We can still do it over time but I would say that would be a separate project. So we don't really know how many would fall back to TCP. We know what some of them say they would by looking at their EDNS(0) numbers. But we don't know that those are true. I have not seen any research done saying, for example, a default configuration of bind submitting an EDNS value of 1000, that it would actually fall back to TCP if you went over that. We just don't know that. So I think that's more research that's not directly involved here but could be spurred out of this. Thanks.

TIM APRIL: Kaveh just lowered his hand. Any other?

DUANE WESSELS: It seems to me like we need to delete that part of the clarification. But then I'm not sure what do we say in response to the direct question—to be decided later?

PAUL HOFFMAN: Or you can simply say, "We removed it." I think that that's fine. What goes to the Board is the end of all of this. You can say, "We decided not to ask." I'm not suggesting we forget about the question. I don't care as much about the first part as I do about the second, about greater TCP, how is that going to affect things. But I think you can just unask the question if you want.

GEOFF HUSTON: I was going to chime in and say unasking the question—sorry. Asking the question is fine. Trying to make it a threshold question is where the problem lies. I actually think there is a lot more to understand around this issue of large responses between the root and recursive resolvers. I've certainly done a lot between authoritatives and recursives and it's certainly some interest of the way they behave to the extent to which there are loss rates and that they do get relatively high even in TCP. There's a whole bunch of folks that just simply give up, which strikes me as being the worst of all possible outcomes. But it's a continuum. And interestingly enough, it's not that the world behaves the same. The world behaves remarkably differently and indeed ISPs behave remarkably differently. So to smack and average number across such a broad spectrum of behavior puts you in the problem of averages that you can see a bunch of detail which is either significant or not significant as a secondary piece of analysis.

I certainly would encourage the work and would also encourage this whole concept of to what extent does a large UDP response tickle a higher TCP behavior rate as being quite significant here. And that's a good thing to do but the results are not "Oh, it's good" or "Oh, it's bad." I would have thought a more meaningful result is, "Well, here's the distribution of behaviors we saw across the broad spectrum of tests. Here is the continuum of sizes and capability." How to assess what's significant or not becomes the next piece of work, but without the first piece of data, the conversation is unanchored. So I would encourage some work on understanding that broad spectrum of behavior, building on existing research work, of which APNIC has certainly contributed a small amount, to understand this problem a little better in the first case

and shy away from thresholds that go black/white. Because I think that's the part that would mislead us all. Thank you.

TIM APRIL: Kaveh?

KAVEH RANJBAR: Thank you. What I remember—actually, all of this work, because RSSAC in the past few years was mainly busy to get to the point where we published 037 and 038 and even after that. And a lot of this work was in anticipation of how can we add new operators, and then we add them, how would you decide how tall you have to be to be able to play or how would you measure performance. This one is not directly related to that. As I remember, it did not conceive as directly as an outcome of that. But the question came a lot of times. How can you add a new one? What would change? And what are the risks and all of that? I think to understand, this is basically an anticipation of the work of GWG and the future of basically root server system, how would you add new operators and what would you monitor? Again, if I remember correctly, that was the reason that we didn't go for any numbers at that point. But of course, if someone must use them as a metric or to check something, then I guess they would have to put some thresholds and see how the system behaves. But my recollection—and on behalf of a few other RSSAC members, you can correct me—is this is an anticipation to have agreed upon measurements, basically, or checks on the future of root server system then you want to add or remove operators.

TIM APRIL: Paul?

PAUL HOFFMAN: Actually, I just saw Duane raise his hand I think to answer Kaveh. Duane, if you're doing that, I would say go first.

DUANE WESSELS: Kaveh, I think it's a little bit different. I think this work all originated with the suggestion to sign the root-servers.net zone. I don't think it was really about adding or removing so much. But the idea was—

KAVEH RANJBAR: You're right.

DUANE WESSELS: We were saying if we were going to sign then why not just rename to make the signing easier. That was really the genesis.

KAVEH RANJBAR: Yeah. This one, you're right. Okay. Thank you.

PAUL HOFFMAN: Going back to the question of do we unask the question? The questions that were pointed out here in the ARR response from us, if the questions remain, would this result in a lot of UDP fragmentation should root server operators expect to see a significant increase in TCP traffic? If RZERC is okay with ICANN answering, "We don't know, here's a bunch of

data that might help inform a future answer," then I don't think you actually need to unask the questions, as long as you're willing to take "We don't know because we can't define significantly and a lot." But here is a foundation for asking those. I think that that would actually be acceptable leaving the questions as they are. If you want real answers then we still have to answer, come up with what you meant by a lot and significant.

DUANE WESSELS: That works for me certainly.

TIM APRIL: We'll just frame it as we will defer to ICANN on what a lot and significant are?

PAUL HOFFMAN: No, no, no, no, no, no. Do not defer to ICANN Org for picking these numbers for you. Bad, bad. I think that it is clear that these questions are real and someone needs to organize a community effort to answer them. I don't know if that would be RZERC starting community effort or RZERC asking RSSAC, especially because the second question of how much is a lot should come from the root server operators, in my mind. But no, it's not that you want ICANN Org to do that. You want that to come from the communities who care about this and who are being affected by it.

PETER KOCH: Thanks. I'm just going to echo what Paul said. I think the one part that affects the root server operators and that is what RSSAC and root ops are for, and that is a stakeholder group that is involved in the process. The other part, the resolvers and the users left behind or less dramatically being impacted is this anonymous unincorporated group of users. I think we have this scoping exercise next on our agenda at some point and I think this is something that I see easily fall into RZERC's scope, not necessarily doing that research but as Paul I think said, organizing the community. Not that I'm eager to do that, but somebody has to. And I'd agree that's better not for ICANN Org to decide, it's for the community and maybe it's our task to find out who that community is or make suggestions how that community can well be represented. Thanks.

TIM APRIL: Paul?

PAUL HOFFMAN: So this may be a career-limiting statement on my part but I completely agree with Peter that that community needs to be corralled and asked about and prodded at even though they really, really don't like being represented and such. These are not the only questions for which we need to talk to a bunch of resolver operators that might even be considered vaguely representative. I can see Geoff's hand waving about that one. But that needs to be done. That needs to be done as a community thing. I agree with Peter. I would say ICANN Org can certainly help on that because that's going to take a bunch of legwork and that's

what we do for RZERC and RSSAC and SSAC and such like that. But I think it needs to be done or else we're not going to get any further forwards on when RZERC does something that affects a number of users, what does that mean? Thanks.

GEOFF HUSTON: I was just thinking through this conversation that there is actually a subtle twist to all of this that is actually unstudied in many respects and certainly needs to be part of any study here. Some years ago in preparation for the KSK rollover, I had the opportunity to compare who was asking root servers versus who was asking authoritative servers. There were certainly some significant deviation in those two sets of IP addresses that there's a community of machines out there that ask the root servers through the DNS protocol incrementally and there's a community of recursive resolvers that we see asking authoritative servers that don't ask the root directly over UDP or TCP.

The assumption, because it's very hard to tell from the work that we were doing, was that they were picking up the root zone through basically the hyperlocal technique. In which case, a lot of these protocol considerations don't directly apply. But we actually don't understand the extent to which that happens, nor do we understand the trends over which that is changing, which actually has a huge impact when you're trying to assess changes to the incremental protocol query of the root versus the zone transfer use of the root and to what extent then that changes that might affect one incremental query have absolutely no impact on the other local serving through its own transfer.

I think anytime you start to embark upon this kind of study around the behavior of incremental query response behavior, once you try to talk about significant or impact, you then have this issue of trying to understand the context in which those incremental queries happen versus other mechanisms that make those resolvers that do hyperlocal unaffected by the changes in behavior of incremental query. So what this is again is saying is it's really hard a priori to put thresholds on any of this or even understand the impact of this if you don't have thresholds, unless you understand the context of the broader way in which recursive resolvers are actually making use of the root and how that's changing. Certainly at the moment as there's more effort put behind, publicizing hyperlocal, I think that the picture is changing. So I think this all summarizes the fact that conclusions are hard to come by in this space right now. But there is certainly value in corralling some effort to understand that behavior in more detail than we understand it today. Thank you.

TIM APRIL: A coherent thought. It seems like we've settled on removing the clarification of what each of these values are that Danielle has already added to the document. I'm trying to figure out whether or not we put some sort of community effort or study of community use and impact that way to create some data that we can then review in [inaudible]. When the RSSAC028 study is done, have a comprehensive look at how this may impact the [inaudible] in users and then make a further recommendation beyond that. Paul has his hand up, now it's down.

PAUL HOFFMAN: Yeah. I put it up when you said try to figure out about the community stuff. I thought you meant as part of this work, which is a little bit difficult because that's not in RSSAC028, then once you said after this. And it could even be concomitant with this but as a separate work item.

TIM APRIL: I was going to try to figure out in my head of whether or not that made sense. I don't think that's something that's added to this recommendation. But it's probably something that we would have to either produce another document or another request to ICANN in some way.

We can proceed with that way for that question. And then next on the list was the R1(c). That's one of the four that we were taking the approach of just regular feedback along with RSSAC. Then R2(b) was the final one for RZERC002, which was just a clarification of the requests to the Org and not to the Board. Are there any other comments or questions on the items for RZERC002? Paul?

PAUL HOFFMAN: I'm sorry. I didn't understand what your answer was on R2(b).

TIM APRIL: I think the text we had was this is a request to the ICANN Org to include risk of operational complexity in the analysis of previous recommendation R2(a).

PAUL HOFFMAN: Got it. Thank you. That's not the way we had interpreted the question. So with that, sure. Adding operational complexity to the matrix that's in 028 seems like a perfectly reasonable thing. Duane, I assume you agree with that?

DUANE WESSELS: I think so. I think I may have written that answer. So yeah, I agree with it.

TIM APRIL: You did write it.

PAUL HOFFMAN: I was worried you were about to say, Duane, "I think I wrote that section in 028," which you didn't. That was that section that got sort of thrown together at the very end. I think that's just fine even though 028 doesn't have that there. ICANN can add it to the matrix. I can't imagine RSSAC saying, "No, don't consider operational complexity." Okay. Very good. That's it then for me on that. Do you need me to stick around for 003? Since that's not my area of subject matter expertise.

DUANNE WESSELS: I think it's up to you. Well, I shouldn't speak. I'm not the chair anymore. I'll let someone else answer that.

TIM APRIL: If it's not your area of expertise, I don't think you would need to stick around for that. We can discuss that—

PAUL HOFFMAN: Are there any open questions? If there are open questions, I can stick around and try to represent us.

TIM APRIL: My next question was I thought there were comments or we had filled up parts of it in a previous document, but I can't remember because I can't find it right now. Duane has—

PAUL HOFFMAN: I'm not seeing any in the document that Danielle shared.

DUANE WESSELS: I would like to have some discussion about some of these. As far as I'm concerned, Paul, you're welcome to stay around. As you said, you're not required to but you're welcome to.

PAUL HOFFMAN: I have a hard stop at the top of the hour but sure.

TIM APRIL: That's fine with me. Duane?

DUANE WESSELS: No. I was just going to say do you want me to dive into what was I going to say about this?

TIM APRIL:

Sure. I think you have the most context on this, even more context on this than I do so go for it.

DUANE WESSELS:

I think there's four or so rows here to look at. I guess I find myself in a slightly awkward position here because as the root zone maintainer representative, it was sort of my plan that a lot of this work would be done by the root zone maintainer, by myself, by Verisign. But the way the recommendations are made to the Board, some of these are a little bit odd I think and could use some clarification. And some of this, there's already notes that Peter and I have been leaving comments and whatnot.

But for example, taking the first one, the first recommendation says the root zone maintainer and root server operators should verify and confirm that the zone maintainer research work, it will not negatively impact and so on. My plan was to certainly make that statement on behalf of the root zone maintainer and in the context of the root zone operators. We have a meeting coming up in July, we would put this on the agenda, and it's my hope that we can get a public statement from the root server operators along these lines.

So that's one example. The second one talks about interacting with the community, making them aware of plans.

PAUL HOFFMAN: Duane, before you go on, just holding our rep, the ICANN representative thing, our response is that ICANN Org would engage with the root zone maintainer and root server operators. From what you just said, that doesn't sound like that ICANN Org needs to engage to do that. It sounds like the root zone maintainer can do their part just without the Board asking them and the root zone maintainer can ask RSSAC. Does that sound—

DUANE WESSELS: Root ops, you mean.

PAUL HOFFMAN: Root whoever. So therefore, ICANN Org doesn't need to do anything on R1.

DUANE WESSELS: Yes. I think that's one way that this could take place. It could also take place the other way where it's also managed more the traditional way where ICANN Org does have its fingers in this part.

PAUL HOFFMAN: That would delay things. Just in that it would have to go to the Board first to have them tell us to do it. I'm not saying we don't want to do it but just be clear which of those two you want. I'm sorry. RZERC should be clear who you want doing it. And I will speak as an individual only, I'm sure we would be fine with the root zone maintainer sort of doing this organizational stuff.

DUANE WESSELS: Okay. Kim's got a comment.

KIM DAVIES: I just wanted to say that I think this is written because RZERC is only chartered to provide advice to the Board. I think there should be an accountability loop that this gets done. If it's something that RZERC feels it's worthwhile recommending then there needs to be some follow through to make sure that it happens, because RZERC can't task the root zone maintainer directly to do this, for example. I think it kind of has to be this way structurally and I think it is appropriate that ICANN have some role to play. Obviously, making it as lightweight as possible, if ICANN is not really doing the work directly would be an objective. And I think the question is, is this the right mechanism that ICANN is putting forth to implement this recommendation? Thanks.

DUANE WESSELS: Kim, with regard to the understanding as presented here, do you think that that's fine the way it is or does it need to be reworded in any way? What would you put in the next column? Yes or no?

KIM DAVIES: I think it's satisfactory if it's satisfactory to this committee. I think what we're trying to avoid is that there was something specific or this didn't go far enough in the view of RZERC so that when we get towards implementation or at the end of this project that there's some difference in opinion, that this wasn't done sufficiently. I think it's really

a sort of a gut check for the members here. Is there any more specificity needed on this or is it sufficient that ICANN Org engage, which kind of means refers this work, I guess, in practice and just monitors that it gets done through the existing accountability mechanism. Particularly with the root zone maintainer, there is a contract there and the formal process by which work gets tossed out and we would implement against that.

DUANE WESSELS: Paul?

PAUL HOFFMAN: It really comes down to Duane is the root zone maintainer volunteering to do this on its own without instruction from ICANN, without having to go through the process. And if not, which I would totally understand, then I would say, like Kim just said, we do need to follow the ICANN Org does this, which means there will be a delay but it will also get done and there will be task lists and RZERC will be informed as things happen. Kim has a good point, which is the RZERC can't tell the root zone maintainer to do anything. But if the root zone maintainer is volunteering to organize this and are RZERC is happy with that, great, but then the recommendations would need to change to take ICANN out of the loop.

DUANE WESSELS: Yeah. That's sort of the gist of it is. I've been working on the idea that the root zone maintainer is volunteering to do these things. But I should

really circle back with my management to make sure that's what they think as well. I don't know who's next in the queue. Is it Tim?

TIM APRIL: I believe it was Kim.

KIM DAVIES: It's an old hand.

TIM APRIL: When I read this originally, I think it was mostly done before I started in RZERC, I had seen that recommendation more as a placeholder for this to be implemented in a waterfall fashion where once one was done, number two would begin assuming that it was reasonable for the RZM and the RSOs to implement it. Then ICANN could start to project the—we're considering doing this outward and then get feedback that way from R2. I think if the root zone maintainer is willing to do this voluntarily, I don't think having the ICANN Org track it—it may impose a little bit of a delay but I don't think it would be too much of a hurdle if it was going to be done voluntarily anyway and if the RSOs were to publish a document saying that they support it as well ahead of time.

DUANE WESSELS: Geoff?

GEOFF HUSTON: I think due process matters in these cases. And due process certainly says that answer reports directly to ICANN Org with proposals and ICANN Org implements that through orchestration of others. I think the understanding is indeed correct as it stands. There's no relationship formally between RZERC and the root zone maintainers and the root server operators per se, and therefore I think we should stick to the formal channels. If others decide to act in other ways, that is entirely beyond our control, and so be it. But I think we need to follow our process correctly and not rely on others to volunteer. So I'm in favor of saying that this understanding is correct, that ICANN Org do engage with the root zone maintainer and the root server operators, because that is what we can do rather than relying on an informal understanding that other folk will volunteer effort, which I don't think is an appropriate process in this context. Thanks.

DUANE WESSELS: Peter?

PETER KOCH: In hindsight, it occurs to me that the phrasing of our recommendations had already taken into account what I think you, Duane, mentioned, in terms of everybody assumed that this or that actor is doing things anyway or is doing things voluntarily. We're phrasing it's probably a precondition for other things to happen without explicitly naming who should be responsible for the results or who should be responsible for initiating that action. What, at least, I did not foresee was this formal exchange and then that now ICANN is understanding or has the

understanding that ICANN Org should do this and that. We don't necessarily get that genie into the bottle for this document but that's probably something to keep in mind that somebody should do something will end up on somebody plate just for the next advice.

In this particular case, I would agree with Geoff that yes, due process matters and we can't provide advice to other bodies rather than the Board. But getting back to the role that RZERC is filling, it's probably not appropriate or not in the spirit of that function to actually now tell ICANN to tell everybody else to do something—ICANN as in ICANN Org in that case—and that is true for the next two items as well, from my perspective. Some things do happen between the RZM and PTI. I'm not sure we need ICANN Org as a catalyst for that on the formal side. It might set a bad precedent if we support this understanding in that strict way. I'm a bit uncertain to how to get out of this.

DUANE WESSELS: I'm glad it's not just me then.

DANIELLE RUTHERFORD: I will just offer in the preamble to the recommendation text, "The RZERC does request that the ICANN Board organize the necessary work outlined in the recommendations," if that's helpful to this conversation. It's highlighted here on the screen.

PETER KOCH: If I just may add, I see my hand open still. I'm not sure that for jumping ahead one line maybe or one line of cells, when we said the DNS and

Internet community should be made aware of plans and so on and so forth, I'm not sure that collectively we had in mind that now ICANN Org go ahead and evangelize or inform people of what is happening. I'm not sure that we had any particular actor in mind. We just said, "Yes, that should happen." And again, that is something that we should be aware for whatever the next advice will be, that this ambiguity will be transformed into some particular task and action for somebody so we might need to rephrase things in the future. And in that case, I wouldn't be against ICANN Org or OCTO in person, so to speak, to do that, but other actors could very well do the same thing. It was a bit down to the question who owns that initiative, which is probably the elephant in the room. Thanks.

GEOFF HUSTON: I'd like to agree with that last sentence of yours, Peter, that what we're really saying, is this would be good for this to happen. We actually think it's so good, it should happen. Right. But the only channel where RZERC can, if you will, make that more formal is through the recommendation going up to the ICANN Board, if I understand the earlier sentence that was highlighted by Danielle. I think it is a bit of a quandary. But our box would be ticked if it happened, almost irrespective of how it happened. But our arm only extends in the direction of pointing a finger at ICANN going, "It'd be good if this happened. You should be responsible, because we can't tell anyone else." But that's not really what we meant. We just said, "It would be good if this happened." So I too am sitting there going, "Well, we can't direct the root zone maintainer or root ops. We can direct ICANN Org but if other folks do it, it becomes a moot point. It just doesn't matter." I sense the ambiguity of this kind of

recommendation that we made where we're saying in almost a passive sense, "This should happen." Without actively saying who should be responsible for ensuring that it happen.

DANIELLE RUTHERFORD: We are now at the top of the hour. I think we've got responses for RZERC002 figured it out and I can put a final copy out to the RZERC mailing list for a final review before I send it back to the ARR team. Tim, do we want to keep RZERC003 on the spreadsheet moving forward and keep the conversation going on the mailing list for how to respond to the 003 understandings?

TIM APRIL: Yes. Do it that way. And if we somehow come to consensus on 003 between now and the next meeting, we can ship it or we can keep the discussion in June.

DANIELLE RUTHERFORD: Okay. Then I think we will close out that conversation for now.

TIM APRIL: We can bring up the two other topics in June. I don't think they're immediately pressing.

DANIELLE RUTHERFORD: Okay.

TIM APRIL: Anything else to discuss? I'm not seeing any hands. Thank you, everyone. Talk to you on June. And thank you again, Paul, for joining.

PAUL HOFFMAN: Sure, no problem. Thanks a lot. Bye.

[END OF TRANSCRIPT]