DANIELLE RUTHERFORD: Hello, all, and welcome to the RZERC May teleconference held on the 19th of May 2020 at 19:00 UTC. Duane, over to you.

DUANE WESSELS: All right. Thanks. Welcome, everyone, to our May meeting. We'll start off with a quick rollcall, please, Danielle?

DANIELLE RUTHERFORD: From the SSAC, Geoff Huston.

GEOFF HUSTON: Hi.

DANIELLE RUTHERFORD: RSSAC, Brad Verd.

BRAD VERD: Hello.

DANIELLE RUTHERFORD: ASO, Carlos Martinez, I note is not on the call. IETF, Jim Reid, I note is not on the call. GNSO Registries Stakeholder Group, Howard Eland.

HOWARD ELAND: Hi.

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: CcNSO, Peter Koch.

PETER KOCH: Yes, I'm here.

DANIELLE RUTHERFORD: Verisign, as the RZM, Duane Wessels.

DUANE WESSELS: Yes.

DANIELLE RUTHERFORD: ICANN Board, Kaveh Ranjbar.

KAVEH RANJBAR: Yes.

DANIELLE RUTHERFORD: And representing PTI, Kim Davies.

KIM DAVIES: Hi.

DANIELLE RUTHERFORD: And then, for ICANN support staff, we have myself, Danielle Rutherford. Duane, over to you.

DUANE WESSELS: Okay. Thanks. Welcome, everyone. So, the agenda you have on the screen, here, and in your e-mail. Anyone have comments or amendments to the agenda at this point? All right. So, we'll proceed.

The first thing to do today is to review or approve the minutes from our February meeting. Prior to the call, Danielle and I were sort of commiserating about why it has been so long since we've met. February was our last meeting.

In March, we had a couple of conflicts, I believe. One of them was my scheduled vacation, which did not happen, and then there was also the IETF meeting. And then, in April, we elected to postpone because, to be honest, there was a lot going on with key signing ceremonies and other things. And so, here we are now in May.

I hope everyone has had a chance to look at the minutes from February and review for corrections. Would anyone like to have any comments or correction on the minutes at this point? Okay. Unless there are objections, then we can take the minutes as approved and Danielle can post them to the website at her convenience.

Okay. Thanks, everyone. So, the meeting today, really, we have a couple of work items to consider. These are proposals for work that RZERC may or may not want to take on. I've sent some PDF files describing these two projects. One is adding some protections to root zone content and the second one signing of the root-servers.net zone.

So, I think you could consider these as sort of draft proposals. Mostly, I'd like to discuss whether or not this work is appropriate for RZERC, whether RZERC should take this on, and if you think these draft proposals are going in the right direction.

Here, this is really the first time RZERC has considered work brought from the committee itself. There have been a couple of times where work has come in externally, but this is the first internal work. Hopefully, we got most of the things right.

Let's start with the first one, introducing root zone data protections. I won't necessarily read the whole thing but the gist of this is that it seems that the root zone continues to get a lot of attention and people are talking about things like hyper-local roots. And we have RFC7706, which has even gone through a second revision.

And the expectation is that the root zone is going to be distributed far and wide, even more so than it currently is. What we'd like to see are some protections to ensure that root zone data that gets distributed and loaded matches the data as it was published by the root zone [19].

Can you scroll down a little bit in the proposal, Danielle? So, the document here just goes into more detail about that; reasons why we think that such data integrity is important. It talks a little bit about how there are some protections today that are of use between the root zone maintainer and the Root Server Operators.

Also, where those fall short in a broader distribution of a root zone. It talks a little bit about DNSSEC and, again, why that is not the ideal solution for what we're trying to achieve, here.

And then, it ends up mentioning there is an Internet Draft, of which I am a co-author. It's called ZONEMD. This is a proposal by myself and others for a new record type that can go into any zone, not just the root zone, for adding cryptographic protections to zone data.

So, the ask of RZERC is to consider both the general problem of adding these protections, and also the specific proposal of ZONEMD. So, with that summary, I'll open it up for discussions or questions if anyone has some. Hi, Jim. Welcome. I see your hand is up, but you're muted, still. Jim, you're still muted. I can't hear you. I'm going to move onto Peter and we'll get back to you, Jim, if you get the microphone figured out. Peter?

PETER KOCH: Yeah. Thanks, Duane. So, first of all, I think it's great that something is brought forward in front of the committee from within the committee and we are no more in desperate need of seeking something to discuss and to shape the tasks of the committee, and so on and so forth.

With the proposal at hand, I have only ephemerally followed the MD record discussion in the IETF, so I claim no knowledge of detail, there. What I'm wondering here is, I understand the hyper-local part but there is also another field of application, which is demonstrating that no records have been removed, or something. Is that a request that is more … Is there any evidence? Has that happened in real life, or is that something that you foresee and want safeguards against?

And then, my second question is, since we're still struggling a bit with the mission, the technical parts of this will be addressed in the IETF and, I believe, the consequences for root operations will definitely be dealt with by RSSAC.

My reading of the charter was that, yeah, we need to look at the problem and whether all the affected parties did have a say. So, what is the exact ask of the committee here? Thank you.

DUANE WESSELS: Okay. So, I'll take a stab at answering those two questions. So for the first, you asked about … Somewhere in here, it talks about detecting additions and removals of root zone data. No, that is not something that has happened.

The reason that I included that is because it's something that RSSAC has been discussing a little bit in the context of this metrics doc that we recently published. I was a co-chair on that work party. In that work party document, it describes root server system metrics that can be performed, and how to do those.

One of those is to ensure that Root Server Operators are serving what they're supposed to be serving. For example, there is a check in there that, if you send a query for something that's supposed to be returned NXDOMAIN, it actually is returned NXDOMAIN, and so on. So, that's kind of where that comes from. Does that answer that question?

PETER KOCH: Yes, partly. Let me rephrase, but take the way you wish. When I asked, “Has this happened?” I didn’t refer to, “Have alterations happened?” The question was, “Were there explicit requests from outside RSSAC to introduce these safeguards?” It’s all level-nine stuff, so on and so forth. That’s just the reason I'm asking.

DUANE WESSELS: Yeah. To that question, I think the answer is no. As far as I'm aware, there have not been asks outside for something like that, exactly. So, the second part of your question was … Can you remind me?

PETER KOCH: If I ever manage to unmute. So, the role of the committee details being discussed in here, in RSSAC, was the exact ask of the committee, as in oversight of all stakeholders being involved, and so on and so forth, or what is it?

DUANE WESSELS: Yeah. So, speaking, I guess, not as your chair so much, here, but as the root zone maintainer representative, the reason that I would bring this to RZERC is because at least the specific proposal of the ZONEMD draft adds a new record type to the zone, and that’s an example of something that we have previously talked about as affecting the content of the zone which is in RZERC’s remit. So, that’s the only reason.

PETER KOCH: Okay. Makes sense to me. Thank you.

DUANE WESSELS: Okay. Thanks, Peter. I'll get back to Jim. Your hand is up, still.

JIM REID: Yeah. Can you hear me now?

DUANE WESSELS: I can.

JIM REID: Excellent. Okay. I think this is certainly something that's within the scope for the committee because of the fact you just mentioned, Duane. We're adding a new resource record type to the root.

However, there are a couple of other meta-issues which I think probably could get sorted out first of all. One of them is, when is this ZONEMD draft ever actually going to merge as an RFC for the IETF? Which I realize is a bit like asking, "How long is a piece of string?"

I think another consideration that needs to be looked at is, how quickly would the existing root zone operators be in a position to implement and support this? I can remember the lengthy discussions, way, way back, about adding v6 capabilities to the root and also getting DNSSEC support at all the Root Server Operators and all the instance they had? And since those discussions, things have moved on quite a lot. For example, I know that the ISC is now using Cloudflare to supplement their existing DNS service with the Cloudflare service support, ZONEMD. I mean, should it be in BIND and [inaudible] NSD and whatever else?

[But we have enough commercial implementation.] Would that impact the way in which the F-Root server's actual service is being delivered? So, I think we need to gauge the level of support or commitment there is going to be from the root zone operators, from a timing point of view, and also from the IETF, before we can get down to looking at this in detail.

And I think another thing that might fall off the back of that is also the question of, how do we maintain some kind of registry or procedure for inserting new resource records and new resource record types into the root zone, and what's going to be the process for that/the procedure for that?

DUANE WESSELS: Okay. Thanks, Jim. So, to the first question about the IETF's schedule, again, as a co-author, in my communications with the working group chairs, they've told me that … So it's come through working group last call, and they've said that this draft is next on their list for IETF last call. You know how this works as well as I do, so take that with a grain of salt. Hopefully, it's going to proceed soon.

Your point about RSO supporting the record is a good one. I think that's something that the committee can ask, or try to find out, or we can try to find out some other way just by pulling the Root Server Operators … If you'd like, we can pull them before the work is started, if you think that's important. Geoff, your hand is up?

GEOFF HUSTON: Yeah. Look, on the whole, I think this is a good idea. The amount of work for the root servers, the RSOs, as far as I can see, the [PE] record, the signed record, the ZONEMD record, is actually part of what's being produced by the authoritative single service and everyone else just serves that as a resource record type.

And so part of the issue is, as Jim pointed out, the adoption of a new RR-type in the root. But to my mind, that is not a major issue insofar as it is largely based on software as it comes out. And it doesn't matter if it's not there. It matters if it's wrong. And so, to some extent, I think the failure point is harmless.

I have one question. Is it going to be signed by DNSSEC? Is it going to be treated like an MX record or is it actually a part of the signed part of the zone, and therefore covered by DNSSEC itself? My assumption is that it would be signed. But I haven't read the draft, I'm sorry, so that was just asking for a point of clarification.

DUANE WESSELS: Yeah. The way that the protocol works is it is authoritative data, so it is signed as all other authoritative records in the zone would be signed.

GEOFF HUSTON: Well, in that case, my kind of response is, let's do it, pending the production of the standard out of the IETF and the agreement with the operators to support the resource record type. But we can't wait in a deadly wait for everyone to make a move, and then we bless it. I actually think we should be saying, "Yes, let's do this," and then leave it to folk who actually are going to implement the fears and devices to then head down that past.

DUANE WESSELS: Okay. Thanks, Geoff. Jim, go ahead. Check your mute button, Jim.

JIM REID: Damn this unmuting thing. Yeah. I kind of agree and disagree with what Geoff said. If we were talking about any other zone, to hell with it. Let's go ahead and just do it. Let's not mess about.

The issue I've got, really, is the fact that we're dealing with the root zone, and for many, many, many people, this is treated as something very special. I don't want us to be in the position of recommending something which other people might consider as being experimental, or that we carry out something which is somehow going to be prejudicial to the operation of the root server system.

Now, that's not the case as far as we're concerned, but I'm sure there are some people out there that are probably going to take that line: "Wait a minute. There isn't an RFC for this? What's going on here?"

DUANE WESSELS: Right. I see what you're saying, Jim. I don't want to put words in Geoff's mouth. He can speak for himself. But I don't think that we would see this deployed prior to an RFC being published, but I think that we can do work in parallel to make it happen on as good a schedule as possible. Was that it, Jim? Okay. Peter, you're next.

PETER KOCH: Yeah. Thank you. Again, please excuse my total ignorance and probably, also, lack of preparation, here. I read the one or two-page summary but didn't go into the Internet Draft, and also not into any discussion. There might be some operational consequences. I assume this is a record that will live at the zone apex and some consideration might be due.

I don't know whether that happens during the discussion in the IETF or would be special for the root zone. But as Jim said, the root zone is always special, or perceived as special, so there is some operational part to be looked at. Also, just because it's an IETF standard, or proposed standard, doesn't mean it has to be deployed instantly.

So, there needed to be a reason and an explanation who the actual user of this is, because the distribution of the root zone between the root zone distribution function and the Root Server Operators is already secured, and that's going side by side. Yeah. But that's not saying it shouldn't happen, just things that need to be on the communication side and should be investigated or researched a bit. Thank you.

DUANE WESSELS: Okay. Thanks, Peter. Howard, go ahead.

HOWARD ELAND: Can you hear me?

DUANE WESSELS: I can hear you now.

HOWARD ELAND: Okay. Thank you. Yeah. So, the point for me is not so much on the actual root zone content because one digest, I don't think is going to hurt anybody, with the appropriate RFC-status caveats behind me.

But to me, as important as pulling RSOs is really going to be pulling what I would call the “validation equivalent.” Right? I'm not exactly sure what that would be called. Because, for example, just because the record is in there, there are folks that want to propagate a root zone of their tailoring. They could simply remove that record and move on, and if the tree is falling in the forest but no one is hearing it then I'm not sure it matters.

So, I think it’s going to depend a lot on the uptake of that from the side of those checking. And then, of course, there’s what happens if your stub resolver, or whoever happens to be hitting this, in a hyper-local scene, whoever is hitting that hyper-local root zone server is going to be trustworthy or not. And maybe trust is a whole other issue.

DUANE WESSELS: Yeah. Thanks, Howard. So, since the zone is signed, and the record is signed, and all that, a consumer of a zone with this record can detect if it was removed upstream, for example. So, that’s one way to address that.

But I take your point about implementation of the validation side. This would need some support in a recursive nameserver software, like BIND, and Unbound, and Knot, and so on. Although, that's not the only way that you can use this record. There are other ways that you can perform validations, but the expectation is that it gets built into recursive nameserver software. Geoff?

GEOFF HUSTON: Yeah, hi. Look, in responding to Howard a bit, there, to my mind the point of this is actually the hyper-local operator. Clients of that recursive resolver who send a query through that can always use DNSSEC to ensure that what they're getting is the genuine article, and this ZONEMD doesn't change that.

But the hyper-local operator may or may not get this directly from a genuine root server and may be confused by a man in the middle and start serving dud data. Now, clients wouldn't be unable to validate that, that's true, but what it would do is effectively negate the utility of that hyper-local service, because it's being given trash and doesn't know it.

So, the issue is, when I get a copy of the root zone and wish to serve it, is that the real deal? And individual DNSSEC queries would help if I, effectively, walk through the entire zone with individual queries, a bit like aggressive NSEC caching, but that's crazy.

This one allows me to get a copy of the root zone without necessarily being sure where I got it from—man in the middle, whatever—and being able to authenticate using the root key that this is a genuine, complete, authentic copy of the root zone.

What that means is that the service that I then push through my recursive part of the resolver-facing clients, to the best of my belief, will have integrity, and I think that's of value. It doesn't solve all the security problems in the world. Of course not.

But it certainly helps the hyper-local operator to ensure that, no matter how they loaded up this copy of the root zone, what they're serving from it, the actual "this is the content," is authentic and real. So, to my mind, that is an incremental improvement on where we are now. It is efficient and it is worth pursuing. Thanks.

DUANE WESSELS: Yeah. Thanks, Geoff. That's a good summary. So, one question I have for the committee is, since we haven't really taken on work like this before, I'd like some input on, what do you see is coming out of this? Do we need … Like, in other groups, we have work parties. Do you see forming a work party around this or writing a short document, a long document? Would RZERC want to do research or reach out to Root Server Operators and ask them these questions that we have about support?

So, I want to gauge the level of work that you think is before us here. Any thoughts? It's okay if you don't know the answer now. We can come back to that on the later call or on the list. Geoff?

GEOFF HUSTON: I think this is more a matter of principle and architecture, rather than mechanics and detail. The IETF is grinding through message digests for DNS zones and one would expect out of them some consideration of all zones, including the root, in terms of the process of making that into an RFC.

One would expect, reasonably, from the Root Server Operators, a relatively exhaustive process of ensuring that they're able to answer requests for a new RR-type, were it to be approved, so that the mechanics of serving it is, again, other people's problems.

The principle and the major issue here, for me, is actually the introduction of a new resource record type into the root zone. And to my mind, it's that sort of single thing that merits the substance of the outcome of this particular group because that, I think, is the evolutionary step.

This document is not far away from that but it probably deserves some further consideration as to the barriers into introducing a new resource record type, the risks and opportunities, and an understanding of, "Well, what if we did it and no one recognized it?" That kind of issue. Analyzing that would be a useful piece of work that we could do. Thanks.

DUANE WESSELS: Okay. Thank you, Geoff. Kim?

KIM DAVIES: Yeah. Thanks, Duane. I guess one thing that's not clear to me is exactly the kind of form that this work would take at the end. Do you have a clear sense of whether it would be in the form of, "Here will be an explicit recommendation to the ICANN Board, adopt this technology"?

Or would it be a set of requirements that, “Should this technology or an approach like this be adopted, here is a set of considerations that should be made in that adoption”? I guess it’s not entirely clear to me—obviously, there are different relationships, here—exactly what form this should take and how that would play into overall adoption.

I mean, for the sake of argument, we could assume that everyone is on board, in principle, with this being implemented as a piece of work, but I'm trying to wrap my head around the order of events and where RZERC’s role would sit within the adoption process from beginning to end. Do you have any thoughts on that?

DUANE WESSELS: Well, mostly I have questions because I have the same questions you do. That’s kind of what I was asking. It’s not clear to me if the output of this is, as you said, maybe a short report with some recommendations that RZERC thinks this is a good idea and that the affected parties should proceed, or is it, maybe, more of a longer document like …

For example, this proposal references work done by RSSAC a couple of years ago. There was a work party that investigated the names of root servers and there was a bunch of work done on setting up fake root zones signed different ways, named different ways, and understanding the size of responses.

So, to me, RZERC doesn't feel like the kind of committee that really does that kind of work, but maybe? I don't know. So, I mostly have questions and not answers at this point. Peter?

PETER KOCH: Yeah. Thanks, Duane. This goes back to my question, what the exact ask is. And again, I would like to offer my understanding. I'd be with you, Duane, when you say that this is not the kind of committee that does this work. I would agree.

It's RZERC's task to make sure that the work that needs to be done is done in the appropriate places, and/or everybody who has a stake in this, or has a say in this, has had the chance to come to the table, and that would include the Root Server Operators.

Of course, the IETF is in the game anyway. And yeah, then, maybe, the other constituencies that send people into the committee, or even people affected outside like the currently unorganized coalition of all the hyper-local operators as but one example.

DUANE WESSELS: All right. Thank you. Jim?

JIM REID: Thanks. I pretty much agree with what Peter and Geoff have said. Perhaps the way forward might be to come up with an idea of what we think would need to be done. So, we could perhaps say, "We're considering this idea of doing ZONEMD for the root, and then send the equivalent of liaisons to all the other necessary affected parties."

Say, "Here's what we understand the steps of this are going to be. Is that your perception of what you think needs to be done? What do you think about the RSSAC, not in terms of the actual proposal, but what your role in this might be?"

Likewise from the Root Server Operators, possibly for IANA, as well. We send [inaudible] himself, and maybe, at some point, also the board, because we certainly are not the ones who can make a final decision about this. Maybe something we have to figure out is who ultimately takes the decision about it if this is to go ahead.

DUANE WESSELS: All right. Thanks. Geoff, go ahead, but I want to wrap this up so we have time to talk about the other topics. So, last comment.

GEOFF HUSTON: Yeah. I very quickly looked at the charter on the web and it seems to me that what Jim and Peter had said actually coordinates, or at least chimes in, with what's on the web. We coordinate, we make sure that folk who are affected are aware, and they have the ability to consider this and make their own decisions.

As to orchestrating this, liaison statements would work with almost everyone except the clients of a hyper-local service. Who is running it? And in some ways, the ZONEMD draft maybe takes on that responsibility, in looking at a section called "Operational Considerations," Duane, that talks about the operational issues around ZONEMD, and how it would help, and what the risks are. And perhaps if one included that and used that IETF/RFC channel as a way of hooking in the operator community, we would be there.

The other way of doing this is to actually do a presentation through something like DNS-OARC, where one gets a reasonable proportion of DNS operators into the room and uses that as a sounding board for the operational part of this, which to my mind is the only bit that’s obviously missing from all affected parties.

So, my incremental suggestion would be to take it to some operational fora, as well, as part of this process of making everyone aware and have the ability to send feedback. Thanks.

DUANE WESSELS: Okay. Thanks. That’s a good suggestion. For what it’s worth, I did present this to OARC but it has been, I don't know, a year and a half. So, it could definitely do with an update.

GEOFF HUSTON: I’ve forgotten you even did it, Duane.

DUANE WESSELS: Well, I know.

GEOFF HUSTON: It was a long a time ago.

DUANE WESSELS: I forgot exactly what meeting it was, and I don't know if you were there. Okay. So, unless there any last-minute comments on this, I want to move onto the other topic in our remaining time. Okay.

The other proposal that I'd like us to consider is signing the root-servers.net zone. So, this is something that has been suggested for a number of years in a number of forums. We've talked about it in RSSAC. I believe it has been mentioned in SSAC, probably more than once.

So, you had a chance to look at this document, hopefully. The gist of the issue or the rationale here is that, although the root zone is signed and although you can form a chain of trust from the root to any leaf note in the DNS that has such a chain, since the root-servers.net zone is not signed there are some potential attack scenarios or traffic-hijacking scenarios to which some people may consider the root to be vulnerable.

For example, if you can convince a recursive name server to cache an incorrect IP address for a root server, that resolver may send some traffic to that wrong root server but you still may get good, valid answers. So, it's essentially kind of an interception attack.

I apologize because I just realized that, when we were previously talking about the ZONEMD, I referenced some other work that RSSAC had done and I had gotten confused. This is the context for that other work that RSSAC had done in understanding the size of priming responses for signed and unsigned zones, and root zones where the root servers had different names.

So, that is certainly one of the considerations of signing the root-servers.net zone. There is an RFC that was published. Sorry, I don't remember the number. An RFC came out a year or two ago that updated the current thinking on priming queries. In that, it says that, when you do a priming query, you should set the "DNSSEC okay" bit so that you get DNSSEC signatures.

And if the root-servers.net zone are to become signed, those priming responses all of a sudden become a lot larger. So, that's one consideration in this proposal. Can you scroll down a little bit, Danielle? Go to the figure. So, the figure just shows how things are today, and the yellow or the gold shows data that's not currently signed but could be signed.

Another thing that, again, I believe was called out in that other RSSAC work, another consideration to signing the root-servers.net zone is that, in order to build the chain of trust and to do validation between a root trust anchor and the root-servers.net data, you have the .net zone in between. So, you need to have a DS-record published in the .net zone.

Some people express concern that that becomes a point for potential failures if the .net zone is unavailable. Or if that DS-record is not correct, then that would be a factor. Some other questions in this proposal get to whether or not current recursive implementations actually validate priming responses.

So, if the root-servers.net zone were to become signed and the responses now have more signatures, what happens if those signatures don't validate? Do any recursive name servers actually check, or do they care? Do they stop working, and so on? Those are semi-unknowns at this point.

Based on some research that I was able to do, it seems like some implementations will validate and others will not. Let's see. I believe that's basically it. So again, I would like to open it up for discussion on whether or not the committee feels this is something that RZERC should take on as work, make recommendations on, or not. Open up for anyone. Go ahead. Geoff.

GEOFF HUSTON: I was looking at the size of the DNSKEY responses to gTLD zones in the root. Over half of them have responses that are more than 1,400 octets. For v6, if the query is over v6, this becomes an amazingly big issue because it forces even the priming query, if you were going to do this right, to go to TCP, because so many folk don't accept fragmented v6.

So, part of the investigation that I think needs to happen here, Duane, is to understand how big these priming responses get and the consequences of hitting over various well-known thresholds—you know, 512, 1,280, 1,500 in the overall response size—and the implications of the robustness of the root because, to my mind, this is the area that we really haven't explored technically.

And just simply saying, "Let's do it," seems to me to be putting it backward. We need to understand what kind of sizes we're dealing with and how those kinds of sizes work for the community of recursive resolvers that perform these priming queries. So, I don't think we're anywhere near that level of understanding.

We're also, I think, unsure how good it is. I notice in the RFC8109. Of which Peter Koch was a co-author, it was a vague thing. The actual text says, “Having DNSSEC validation for the priming queries ‘might’ be valuable.” Not “will,” “might.” And I think, maybe, that needs to be explored, as well. Thank you.

DUANE WESSELS: So, Geoff, do you think that RZERC itself should take on this work, or ask some other body to do that?

GEOFF HUSTON: Well, from the last discussion, I think it’s a coordination of other people doing some work here.

DUANE WESSELS: Okay.

GEOFF HUSTON: There is some work from the measurement community. There is certainly some work that [mind we are] the Root Server Operators or DNS measurement/DNS-OARC style work. But this certainly is a topic worthy of study, and I am not sure that just this set of individuals here is sufficiently armed and ready to perform such a study. I'm sure, if we hoisted the flag up and said we’d be interested in understanding this, researchers, and measurement folk, and so on, might give us a hand.

DUANE WESSELS: Okay. Thank you. Peter?

PETER KOCH: Yeah. Thanks, Duane. So, when the work on the priming RFC, or the draft at the time, started, we quickly came to this question of signing that zone, and part of the discussion that is reflected or reoccurred, being quoted in the paper you submitted, was performed, but rather inconclusive.

So, the root servers serve the root and root-servers.net, but obviously not .net, and what's going to happen, and how does that benefit the consumer? So, yeah, more study is needed. Maybe measurements, but actually, maybe, active measurements in lab environments. I think that is what Geoff probably meant when he said "studies to be performed," and not by this committee, I would completely agree.

On the other hand, this seems to be very much related to recent suggestions that, "Oh, if we only had decided to sign parent-side NS-records, signing the delegation to there is even a bit of protection if the child isn't signed," and that means we are going back to … Maybe not square one, but square two, or something, in the DNSSEC discussion and the architecture of DNSSEC, just that we are doing it at a very special place in the tree.

So, that's another discussion that needs to be had. It would probably not be necessary to do this if the deployment status of DNSSEC was higher because, as was mentioned, and I think also mentioned in the paper, this is always only error detection, and not error avoidance. So, if NS-records are forged in the delegation—or in the referral response, I should say—then DNSSEC will do its job.

Now, if the majority of children aren't signed, then DNSSEC can't do its job, and here we go. So, I'm wondering what the architectural implication is and why, again, the root is so special. This is an interesting discussion to be had, probably with the IETF, but maybe not only in there. So again, it would mean that there needs to be an architectural discussion, cross-community, as we like to say in some places, with the operator community and those parts who are responsible for deployment to get to the real problem. Thank you.

DUANE WESSELS: Okay, thank you. Jim, go ahead.

JIM REID: Thanks, Duane. I pretty much, again, agree with what Peter and Geoff already said. To my mind, I think this is certainly something that RZERC can tackle, at least in terms of scoping things, and perhaps defining what are the known-unknowns, and who could maybe look into them.

Obviously, there will be a need for measurements. Probably a vote from RSSAC, the RSOs, as well. So, I think it's certainly worth us taking on some piece of work here but, again, not actually to do the detailed measurements and metrics that might be part of this exercise.

And just throwing something out there for the sake of it, why did we have root-servers.net? Why can't we just put the address records/MX-records for the root servers in the root zone itself?

DUANE WESSELS: So, yeah, that was the whole focus of this RSSAC Work Party from a couple of years ago. I believe its output was RSSAC028, as you can sort of see at the bottom of the screen. That work did investigate different ways of naming root servers, including directly in the zone itself. It studied, how do things look in terms of response sizes and other things when you do that?

For whatever reason, that RSSAC document recommended the status quo. At that time, it did not recommend making any changes to the names of root servers, but I think that's something that people always like to talk about. Howard?

HOWARD ELAND: Yeah, thanks. Jim stole a little bit of my thunder, there. I was going to kind of go along the same lines with that. So, maybe the purpose of this group and put it in a more generic context. I would say, maybe, the thing to gander at is the prospect of signing any delegated zone by which the root depends.

So, whether it's under .net, or somebody did something akin to what they did to the new TLD round, where they threw out anything under .nic, or if it was right in the zone itself, perhaps the way to phrase it is to just look at those different scenarios in the generic form and see how those dependencies lay out in terms of sizing and what that means in addition to some of these other potential gotchas.

DUANE WESSELS: All right. Thank you, Howard. Jim, is your hand up again, or did you still want to comment? No? Okay. Geoff?

GEOFF HUSTON: I agree, such a study should be done. In an ideal world, or one version thereof, we would have a budget. We would be able to commission such a study and pay for it. In the world we live in, volunteers have to do this work on their own coin, for their own benefit.

We can ask that the work be done. In theory, we could do the work, but we, too, are just a bunch of volunteers. And to some extent, I'm not sure that that creates a credible answer for something with the gravity and role of the root.

So, yes, there is an area we don't know about in the DNS. That's true. And I would have thought, in terms of how long our arms are and what we can affect, I think we can certainly flag this as an area of interest and should have some investigation.

Particularly, should these names be signed? And in the validation process, where is an optimal place to have those names reside? Because I suspect it's part of the larger package that Howard alluded to that is not only about signing the names at the root servers, but what are those names?

So to my mind, I'm not sure that this is work that this particular committee can do from start to finish. In fact, I'm pretty convinced it's not. But in some ways, trying to understand who would have a stake in this, and how the work would happen, and who would do the work is part of our process of cogitation.

We need to understand how to see that work. I think it’s good work and we should pursue it further, or get others to pursue it further, but exactly how is something that still eludes me. Thanks.

DUANE WESSELS: Yeah. To your points about us not having or having budget, certainly we don’t have any budget that I'm aware of, but I guess that doesn't mean we can’t ask. I mean, we could consider a recommendation, I guess, to the board to fund such a study, with RZERC, perhaps, as oversight. Is that something that would be interesting to committee members to pursue?

GEOFF HUSTON: I'm not sure we want to morph into being the contracting body here, Duane. That is also a lot of work. I’ve been in other places in the ICANN community where that has happened. It has its own issues. It is an option, that’s certainly true, Duane. But I'm not sure that it would be a preferred option.

DUANE WESSELS: So, maybe not taking it as far as us having oversight, but requesting funding of a study and making oversight someone else’s problem. Jim?

JIM REID: Yeah. A couple of points about that. I think the idea of commissioning a study might be a little bit premature. I think we probably need to do a little bit more work trying to scope the extent of what’s going to be involved.

That scoping exercise might well say, “Well, there’s so much stuff that needs to be done,” once we’ve actually itemized this stuff, and who all the various different component parts this is going to potentially touch. We might reach the conclusion that, yes, this does require an in-depth study, and that needs to be paid for by somebody.

And following from what Geoff said, I strongly agree that we should not, in any way, shape, or form, be involved in being either contracted party or having oversight of the contract. And speaking from personal experience, I’d be very, very reluctant for RZERC to be involved with ICANN’s contract people at all. I think that’s just going to be a very, very unpleasant experience.

DUANE WESSELS: Okay. Thanks, Jim. Peter?

PETER KOCH: Yeah. I'm probably in-line with the previous contributors. I think it’s too early to ask for the money, and even if we do, we don’t want the people to spend it, or the people who decide where it’s going, as in contract oversight or anything.

But it’s also not necessarily clear to me that the money even comes from ICANN. So, I think we should take this up as a work item in one way or another. The first step would be framing the question, or phrasing the questions.

And maybe we can have some informal conversations with potential sponsors of such studies so that the committee, in the end, just has to say, “Somebody should do something.” And of course, the question is, “Okay, are the researchers biased or are they not?” But it’s not clear to me that we want to appear as another ICANN-money-burning committee. Thanks.

DUANE WESSELS: Okay. That’s very good feedback. Thanks, everyone. So, I think we have probably exhausted that topic. Thanks for the feedback. Let’s go back to the agenda. I believe the only other thing on our agenda today was any other business.

Does anyone have any things they’d like to bring before the committee, mention goings-on, or anything like that? All right. It doesn't seem like it. So, our next scheduled meeting will be … Let’s see. We have a meeting scheduled in June. It would be June 16th, I believe.

DANIELLE RUTHERFORD: That is correct.

DUANE WESSELS: Which is a week before the virtual Kuala Lumpur meeting. I mean, it sounds to me like we have topics to continue discussing. Is everyone okay to meet in June, one week before ICANN? All right. Speak now, or forever hold your peace, or bring it up on the mailing lists. If not, we will plan on having our June meeting. Thank you, everyone, for your time today. I appreciate it. We’ll see you online.

[END OF TRANSCRIPTION]