DANIELLE RUTHERFORD: Hello, everyone, and welcome to the RZERC Monthly Teleconference held on Tuesday, the 15th of February 2022 at 19:00 UTC. Tim, would you like me to start the roll call?

TIM APRIL: Yes, please.

DANIELLE RUTHERFORD: Tim April, IETF?

TIM APRIL: Present.

DANIELLE RUTHERFORD: Carlos Martinez, representing the ASO, I note is not on the call yet. Peter Koch, ccNSO?

PETER KOCH: Yes, present.

DANIELLE RUTHERFORD: Kaveh Ranjbar, ICANN Board?

KAVEH RANJBAR: Present.

DANIELLE RUTHERFORD: Kim Davies, PTI?

KIM DAVIES: Present.

DANIELLE RUTHERFORD: Howard Eland, RySG?

HOWARD ELAND: Yes, ma'am.

DANIELLE RUTHERFORD: Daniel Migault, RSSAC?

DANIEL MIGAULT: Yeah, I'm here.

DANIELLE RUTHERFORD: Geoff Huston, SSAC?

GEOFF HUSTON: Good day.

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

DUANE WESSELS: Yes, Duane is here.

DANIELLE RUTHERFORD: All right, Tim, over to you.

TIM APRIL: Thanks, Danielle. So if we could pull up any adjustments to the agenda or switch over to the minutes from last month. Does anyone have comments or adjustments to the minutes? I'll note that there were a couple of actions that due to unforeseen circumstances didn't get taken care of in integrating month. Okay. I can take that as approved. The first item is up to you, Duane.

DUANE WESSELS: I'm happy to go over the draft ZONEMD deployment plan today. Thanks for putting it on the screen. I doubt this will take a lot of time. But I'll just sort of go through it bit by bit. I'm happy to take questions as we go.

Can you scroll down past the Table of Contents? Okay. So, the background, I suspect this is material that everyone here already knows, more or less. It references the RFC and briefly talks about how the ZONEMD protocol works, essentially conveying a digest of the zone data inside the zone itself, how that it can be used with DNSSEC for strong authenticity guarantees.

The next paragraph talks about how the zone is currently distributed among the 1400 or so separate instances, and then makes reference to RFC 8806, which is the IETF's version of hyperlocal root.

And then there's this figure which shows sort of the flow of root zone data through the system and shows the point at which where the ZONEMD record is sort of inserted along with the signing process, and then how it flows the different other parties. So this figure sort of backs up the following section, which talks about impacted components and parties. This is a figure that has been shown in other contexts before or something similar.

One thing that's different here is that this one includes the internic.net services, one of which is currently operated by Verisign, and then the other two are operated by ICANN down at the bottom there. Those are sort of shown as a way that somebody who wants to do RFC 8806, that's a location where they could get root zone data. There's also a line that shows that they can get it from some of the root zone operators. So the bottom sort of represents that aspect of this, which is the resolvers doing their own local serving of the root zone, and the top represents the traditional root server system with queries and responses. Any comments or questions about the figure before we go on? Okay.

Like I said, these next sort of subsections just go through this one by one and briefly say whether or not a component or party is impacted. Since the TLD managers and the IANA RZMs are sort of upstream of where the ZONEMD record is added, there's no expected impact to them for adding the ZONEMD record. It really sort of begins next in Verisign's

root zone management system and distribution system. The Verisign RZMS, the work is already under way to support ZONEMD in those components. And we have verified that the distribution system, those are operating correctly on zones that include the new record.

Then moving along the top in the figure, the root server operators have also been asked to ensure that their internal distribution systems and name servers can support the ZONEMD record. I have some of that feedback already but some is still to come. And it notes here that the root server operators are not being asked or required to enable ZONEMD verification but simply to confirm that the presence of the record does not negatively impact their services.

The next one is the traditional recursive resolvers. There's no expected impact here. We have no reason to expect that those resolvers would issue queries for the ZONEMD record, and if they did, there's no harm that they shouldn't do. That's not part of the verification step, and so they can do that if they like.

The first place where we get to components that may have an impact is recursive resolvers that implement 8806 or other ways of serving root zone data locally. There's not a big expectation of impact because the software versions that support these features are already known to handle the ZONEMD record. That's what it's there for. So those implementations are well poised to work with the ZONEMD record in the zone.

So moving down to the InterNIC-related services. So this refers to the web servers and the FTP servers that publish the root zone and other

zones to the public. There's no expected impacts here because they just treat these as a big file, essentially, and don't really care too much about the contents of those files. But the consumers of data from those services, so users of the InterNIC services, may be downloading the zone from there and running their own processing on the zone perhaps with custom scripts. If those scripts and tools that read the root zone are unaware of the ZONEMD record, their processing may fail. So that's sort of the group that at least I'm most worried about in terms of adding the ZONEMD record.

We'll see later on how we're sort of addressing that in a couple of ways. So any questions about impacted parties before I continue on? Okay. Oh, I see a hand from Peter.

PETER KOCH: Thanks, Duane. Just a minor question. You've mentioned BIND and Unbound in particular. Are there plans to give other vendors an opportunity to make statements, or will these detailed version numbers disappear from the document? What was the plan there?

DUANE WESSELS: Good question. At the time that this document was written, Unbound was really the only, I'd say, implementation of verification. We have been working with some other vendors to get the protocol implemented in their products. So that's PowerDNS and Not Resolver. So since this draft was initially written, that work is now more complete. I don't think those versions have actually been released yet but they're sort of scheduled to be released soon. And we could add those version

numbers if you think that would be useful. I don't have a problem with that. I would have to do just a little bit more research to understand whether or not those implementations are enabling it by default or if it has to be configured. I know in the case of Unbound, it will work by default. If it finds the ZONEMD record, it will do the verification by default.

PETER KOCH: Thanks, Duane. Understood. I think it might not be necessary to strive for a complete list of all the minor versions, and so on and so forth. I wanted to understand what the plan is in terms of using these just as examples or providing an exhaustive list. I have no preference either way.

DUANE WESSELS: If other people have thoughts on this, please speak up and let me know. I'm happy to revise the document to account for that.

All right, so I'll move on to the next section, which is some of the operational details of this. So if you don't know, the ZONEMD RFC defines two standard hash algorithms, SHA-384 and SHA-512. Both of these are sort of, I would say, modern hash functions. The 384 produces a digest of 48 octets, and SHA-512 produces 64 octet digests. And then there's a number of algorithm code points reserved for private use.

So as stated here in bold, for compatibility reasons, we recommend that the root zone use the SHA-384 hash algorithm initially for the reason stated below. Before ZONEMD was an RFC, when it was still a sequence

of Internet drafts, there was an implementation done in BIND. And at that time, when the BIND implementation started, the ZONEMD record, it had sort of the same fields but the fields have different meanings. And initially, the digest algorithm, the hash algorithm was in the second field. And later on in the progress of the Internet drafts, that actually moved to the third field. Both of those fields were eight-bit values so the record format didn't change. But the hash algorithm moved from the second field to the third field. Versions of BIND from that initial implementation only expected SHA-384 digests and therefore always expect that the digest field itself will contain 48 octets. And they will produce errors for ZONEMD record that does not contain 48 octets.

So that was corrected in April 2021 and released in the BIND versions listed here. But these newer BIND versions are not as widely deployed as the older ones. And so there is a fair amount of deployment of BIND that would not be able to accept a SHA-512 digest. They would produce errors. So I think that's a pretty clear choice/reason for going with SHA-384 initially. Ideally, it would be nicer to go with a longer digest.

DANIEL MIGAULT: Just a question, if it fails, does it mean that the versions are not up to date? Because then it might be a good reason to move to SHA-512.

DUANE WESSELS: Daniel, you're sort of suggesting that if the root zone use the SHA-512 digest, that would encourage people to update their versions of buying software? I think in this case, that would not be a great idea. You would find that there's a lot of this slightly older versions of BIND still in

standard operating system installs. It's not really that old. It's only maybe a year old, not like 5-10 years old. So I think that would not be a really great approach in this case. Does that make sense, Daniel?

DANIEL MIGAULT: Sure, sure.

DUANE WESSELS: Okay. So let's scroll down a little bit to the next section. So another choice to be made is about the presentation format. So when the root zone is being distributed from the root zone maintainer to the root zone operators and all the instances, that all happens using DNS zone transfer in wire format, and there's no concerns there about presentation from it. Where it does become a concern is for the FTP and web servers, those files are obviously presentation format for the zone. The concern here is that people who take the files from those servers run their own processing on them, they may not be prepared to see a ZONEMD record in the format as shown here with the RR type set to ZONEMD. If they're only expecting to see NSA AAAA DS records, then their process, it could fail on this ZONEMD record.

We think it's much more likely that the older software would accept the ZONEMD record in the generic format defined by RFC 3597. And if you scroll down a little bit, you can see what that looks like. So those two records are equivalent. But in the generic format, it doesn't say ZONEMD. It says type 63, and then this long, essentially raw string of data. So out of some caution, we're proposing that in those, when it's in

presentation format, the ZONEMD record should be given in the generic format as shown here.

The next section talks about a phased approach in which we're proposing that for the first phase, the record would be published using a private use hash algorithm number. So this is very similar to what was done for DNSSEC in the root zone many years ago, what we called the deliberately unvalidatable root zone. This is essentially unvalidatable ZONEMD record with the private use hash algorithm. It means that the data itself can't really be verified, and we can be sure that only the presence of the ZONEMD record doesn't cause any problems. And then in the second phase, the hash algorithm would be changed to the SHA-384 algorithm, and then the record becomes verifiable. We expect that the first phase would last for about two months. And then after that, the second phase would begin.

A note on correctness checks. Already today there's a lot of correctness checks built into publishing a root zone. But for this, we're adding two ZONEMD verification checks with code written by different authors and different programming languages. And if those verification checks fail, just like any other checks might fail, then the candidate root zone goes into a holding state and will not be published until someone from Verisign would manually check the zone and take corrective action if necessary.

All right. So next is the deployment schedule. We already know some of this that obviously Verisign needs to complete its implementation and deployment of its software that supports this. That work is expected to be complete by Q2 of this year. And then, as I said before, we're still

waiting for some of the root server operators to confirm their readiness for the ZONEMD record. This talks about the two months for phase one and the option to back out, roll back, if necessary to revert to publishing the zone without a ZONEMD record, and then the second phase using the SHA-384 hash algorithm. Again, if problems are encountered, then we have the option to go back to publishing the zone without the ZONEMD record. And that's at the end. I welcome any more questions or comments from anyone. Geoff?

GEOFF HUSTON: Look, Duane, I've just been comparing this draft to RFC 8976. And there are a couple of questions in my head, one of which, while you're talking, I was able to answer and the other one I haven't checked. You advocate using a presentation format which is—I've forgotten the word used, generic whatever, it's format, right?

DUANE WESSELS: Yeah.

GEOFF HUSTON: Whereas 2.3 of that RFC gives the presentation format, which is, if you will, the real one as distinct from type 63. There's no mention in the RFC of a presentation format that varies in any way. I understand what you're trying to do. And you're saying there may be folks who will see the root zone but are not aware of 8976 and were going to use a different format. But you haven't touched upon what part of the

standards digests like this RFC permits that and why it's a variance with the RFC.

The other one that I haven't—because you stopped a bit early before I could read through it and listen to you at the same time—is what happens when you get an unknown SHA digest and how does that correspond to Section 4, which talks about verifying the zone digest? While where root zone key issues had a defined behavior, if you didn't recognize the algorithm, I am trying to understand what the RFC says if the digest is not recognizable or unprocessable. It strikes me that the issue is if the digest is meant to protect people [marking] with the zone, if the digest itself seems corrupted, the action to ignore the digest, I would have thought, should have been in the RFC. And I haven't had the time yet to check that but that default action seems security-wise to be a less than optimal. If I get something that I don't recognize in a ZONEMD, I really wonder where I should keep hanging on to that zone file.

So I suppose the basic question is, have you checked the compliance of this document against the RFC and where does it vary? And secondly, if it does vary, do you intend to change your draft or change the RFC? What's the way through this? Because I'm a little bit concerned about pushing forward a document that doesn't conform to this IETF standard. Thanks.

DUANE WESSELS: So to your second question, Geoff, I guess the part that you ran out of time looking for—

GEOFF HUSTON: I did.

DUANE WESSELS: Yeah. I just brought it up. The RFC talks about this in Section 4, item five, sub item C, and it says, "The hash algorithm field must be checked. If the verifier does not support the given hash algorithm, verification must not be considered successful with this ZONEMD RR." So that just means that it can't use that particular record to verify the zone. If that's the only one, then the verifier has no way to verify the zone. If there are multiple records with different hash algorithms then it can proceed to using a different record or different hash algorithm to perform the verification. So it's sort of similar to the way it works with DNSSEC and those DNSSEC algorithms.

DANIEL MIGAULT: Unsuccessful does not mean fail in that.

GEOFF HUSTON: Well, that's the bit that I can't find. Where does unsuccessful mean don't fail? Because the DNSSEC, for example, an unknown algorithm is explicitly said you must treat this as unsigned. But here, the draft says you must treat verification as an unsuccessful outcome and you should report the reason. I know we're going sideways here but it's a useful sort of check. I can't see clearly that the RFC with all of these authors. It didn't actually say what unsuccessful meant. Or it's written somewhere

that I can't find it because it is a big document. There's a lot of stuff in there.

DUANE WESSELS: Yeah. That was intentional in this document that the RFC doesn't say that a zone should be prevented from loading a zone because the ZONEMD verification fails. So it's sort of implementation dependent, I would say.

GEOFF HUSTON: I don't know if that's the appropriate advice from a security consideration or not, and maybe that is by the by. But in some point, I think maybe in this deployment plan, you should note that in this first phase, it will cause an unsuccessful attempt to verify as per the RFC and implementations are expected to accept the zone in such a case. And that's, as I said, not stated in the RFC.

DUANE WESSELS: I get your point. Yes. I'm happy to add that. Do you want to go back to your first point?

GEOFF HUSTON: Yes. The textual representation format where the RFC says it shall be written in this format, 2.3 ZONEMD presentation format. And there's no reference to the type 63 format.

DUANE WESSELS: In this case, I would say that the generic format in any record can be presented that way, right? It's just an encoding of sort of the raw R data of the record. So this is not something that's unique to ZONEMD. You can write a zone that presents A records, NS records, anything in this generic format. So I guess that's why we didn't feel the need to specifically put a type 63 example in the RFC.

GEOFF HUSTON: But what the RFC did say is this is the presentation format and what this draft deployment plan is saying is, "We're not going to use that." That's sort of what it seems to me. Now, that alternate presentation format that is being advocated in this draft deployment plan is, as you say, written up in other RFCs is what do you do when you don't have a presentation format for that RR type, and that's true. It just, as I said, struck me as slightly anomalous that the RFC says one thing and this deployment plan document is relying on, if you will, different RFCs to say something else.

Again, I can understand the reasons you're trying to sort of do this failsafe and making sure that even if you don't recognize it, you're not going to throw a wobbly on it and reject everything. That's fair enough. But maybe you either reference an RFC that gives this generic format or something. But like I said, a very precise reading of the RFC says you can't do that, the presentation format in 2.3.

DUANE WESSELS: So the deployment plan, it does reference that, whatever it is, 3597, the one that describes the generic format. I guess the way I would think of it

is that if you're a ZONEMD aware implementation, then you should use the ZONEMD presentation format. If you're unaware, then you should use the generic format. Obviously, that's why we're doing it out of the abundance of caution. I'll let some of the others go. Peter, I think your hand was up next.

PETER KOCH: Thanks, Duane. I've just pasted a line from RFC 3597 into the chat that actually explicitly allows the use of this generic format from the implementation side. So I'm not sure I completely understand what Geoff's after here. It might add to the confusion. First, I should say I like this very careful and staggered approach. What might be missing is defining a criterion for changing this code back to the generic format or define some prerequisites for actually moving to the generic format.

A warning for the reader. Yeah, it might be useful but I don't think that by a presentation format being defined in RFC 89-something, it is precluded to use this generic format, especially because 3597 seems to allow this in particular. But then again, if it's to minimize confusion, that particular paragraph from 3597 might be referenced if that addresses just concerns. Thanks.

DUANE WESSELS: Thanks, Peter. I don't see anything in the chat, by the way. I don't know if your paste is still pending or something. Oh, there it is. Okay.

PETER KOCH: Sorry. I didn't hit Return.

DUANE WESSELS: All right. Thanks. Tim?

GEOFF HUSTON: Very quickly in response, if you may.

DUANE WESSELS: Go ahead, Geoff, sorry. Go ahead.

GEOFF HUSTON: If you added even that as a footnote, I think that that covers it off. It seemed anomalous without it. That's all. Thank you.

DUANE WESSELS: Okay. I'll do that, Geoff. Thanks.

DANIEL MIGAULT: I was going to say the same thing of having that as a footnote, just so that people understand that—the default presentation format is the type 63 format unless it's overridden by the RFC where you can change, put the RR, type in directly.

The thing I was going to follow up on those, I think the two-month phasing allows for either OCTO or discussions at IETF or ICANN meeting so that people can be aware that this is happening. I wonder whether there was criteria for a rollback from either phase one or phase two of if there's significant impact noticed. How should someone raise that

concern, or if there's a criteria for a rollback, or you either need to fix it or you're going to be operationally in trouble?

DUANE WESSELS: I remember having lots of discussions about this for KSK rollover stuff as well. I don't have a good answer for what are the impacts. I think we're going to have to sort of just watch it closely around these dates and see what happens. But I do take your point that there's no defined communication mechanism or where people can send problem reports to. So I can definitely work on adding that, if you think that would be good.

DANIEL MIGAULT: Yeah. I feel like at least a notification to DNS [SOC] or something like that. I don't know where else it should go if this change is happening on this date.

DUANE WESSELS: Yeah. I'm definitely planning to do that. But I will make that more explicit and maybe document that. That would be a good place to watch for announcements. Then we will obviously be watching that form as well. Any other comments?

GEOFF HUSTON: I have one. Again, I haven't read the RFC thoroughly. But when you talk about root server operators on page three, when you say they're not required to enable zone in the verification but simply publish as is, you

don't cover the case that they do, and for reasons—cosmic rays, whatever, global uncertainty, Heisenberg—it fails, what happens then? Should you simply say they should publish as is and not check? Or if it fails verification, they should question mark? The problem is whenever you get sort of two conflicting instructions, authoritative server, the master says publish X, and now the ZONEMD says, "Well, X is bad," do you publish, do you not publish? Or are you willing to say this will never ever happen? Whatever you do with verification, it will never be inconsistent, which seems optimistic in the larger scheme of universal uncertainty.

DUANE WESSELS: Yes. In all of the discussions that I've had with the root server operators about this, none of them are planning to turn on the verification in the near future. They're all taking a cautious approach to wait and see how this goes. They'll do some tests. Maybe they'll turn it on. It's as like a warning level, right, and see how that goes. It feels to me like that is pretty far down the road still. I think it's an important discussion to have. I'm not sure that it belongs in this document.

GEOFF HUSTON: I think it is. Because in some ways, when you say RSOs are not required to enable ZONEMD verification, the document then sort of leaves that blank as to what if they do and what if it fails? And while the next sentence actually is back to don't do verification, or if you do, do not act upon its outcome, you should be confirming that the presence doesn't negatively impact operations, full stop. That's, I think, better advice

than the dangling uncertainty of, "Well, the RSO is meant to invent their own rules if it fails, you know." I don't feel that this is solid. That's all. Sorry. I'm being very pernickety today, but maybe this document actually does need that level of detail.

DUANE WESSELS: Sure. That's why we're here. I can definitely see where maybe the document should say RSOs are not required to enable verification at this time. That requirement could change in the future. I do think that you're going to see a lot of RSOs waiving their independence flag and they'll each come up with their own approaches to handling a verification failure. Obviously, in my opinion, it if there is a verification failure of ZONEMD, it's better to continue serving the old zone for at least a short while until the verification failure can be figured out, what went wrong there. But I don't know.

GEOFF HUSTON: You know, Duane, I think that's a big issue, not an RSO independence issue. I think that transcends operational issues of being an RSO. That's about the semantic behavior of the root zone and actually merits deeper discussion at some point somewhere in this ecosystem to actually define consistently behaviors in that space. So I would not necessarily say, well, it's up to each RSO. I think that would be poor advice to the ecosystem at large. That's why I sort of bring it up. I would actually put in title words right now to say they're not required to enable ZONEMD verification. And the outcome of that verification should not affect the publication of the root zone at this point in time.

Rather than simply confirming that the mere presence does not negatively impact then follows logically from that.

DUANE WESSELS: So you're suggesting that the deployment plan document should have language or requirement that says RSOs must not fail to publish the root zone due to a verification failure at this time?

GEOFF HUSTON: Correct. Because that is actually consistent with confirming the mere presence does not negatively impact their operation. So yes, it's taking a very soft line and leaving it undefined in the future as to what should they do if it fails? When we're happy that the presence doesn't cause a problem, I think it's actually a better way of doing this, rather than leaving it unstated. But others might have a view. So I'll stop here.

DUANE WESSELS: Yeah. Peter?

PETER KOCH: Thanks. I think that paragraph has another ambiguity because it isn't clear to me when it says RSOs are not required to enable ZONEMD verification, is this normative or descriptive? I would not know how it could be normative in this place. So for clarity, it would benefit from an explanation why that is the case. Or maybe make clear that it is not this document that says they are not required. This non-requirement originates from elsewhere.

Also, I see that there might be in symmetry because nobody—at least the others didn't so far—makes any statement how the root zone on its journey from the root zone maintainer to the root server operators, and further down to the root server pieces of software, how the integrity is checked. This adds a non-requirement. Again, it doesn't say where that non-requirements—how that is justified. So I understand that it is helpful to manage expectations in a way, but it's unclear who is in charge of making these statements with what strength.

DUANE WESSELS: Thanks, Peter. So I just want to make sure I understood. Your concern was only about the current per wording around RSOs are not required to enable verification? Was your second part talking about a different part of the document, or is it just that one?

PETER KOCH: Sorry. Same part of the document. When it says that they are not required, I'm missing the part because what? Because this document says so?

DUANE WESSELS: Okay. All right. Thanks. Tim?

TIM APRIL: I was listening to the conversation a couple minutes ago. I was wondering, is there a defined escalation path in case that someone notices an error with ZONEMD? Are their verifications broken? Or they

can show that they're doing verification properly and it's something on your side or the Verisign side or is it my default channel of send mail to you and Kim and/or DNS [SOC].

DUANE WESSELS: I think this is something we've talked about already about, that the document doesn't have anything like that, where to send complaints, if you will. So I can figure out what to put in there, some e-mail alias, maybe one specific for this purpose that will go to Verisign and IANA for notifications. Is that kind of what you're getting at?

TIM APRIL: Yeah. I was just trying to think of what my inbox would look like if I tried to do something like this. It didn't work right.

DUANE WESSELS: I will add that for the next version, I guess, and maybe work with Kim if we want to do some kind of joint alias or something like that, I don't know.

With that, at the question if the RSO is currently validate the root zone [inaudible] before publication? I don't know. To be honest, I don't know for sure. If they do, it's something that they've taken on themselves to do. I have heard in the past of root server operators that did do some sort of validity checks on the root zone before they published it. But I don't know how many of those are still in place, to be honest. Kim, your hand's up.

KIM DAVIES: Yeah. I just posed the question to illustrate that I think some of the things that have been brought up today about ZONEMD actually are good questions but also in a broader context. Is any kind of validation done with DNSSEC already that might gate publication of the root zone? It's probably an area to get a better understanding of any way in the broader context to inform anything that might be specific to ZONEMD.

Actually, that last question about the escalation path, same thing. If someone noticed some issue with validating the root zone, I'm not sure we have necessarily a well-defined escalation path, that kind of issue that might be identified that somehow, some way, the root zone has a problem with it, whether it's ZONEMD or something else. But I think that's necessarily well-defined. We do have a root management generic e-mail address that people use for those kinds of inquiries but that's kind of I think just by default. So possibly a good evolution here is—I'm sure Duane and I will discuss it moving forward is beyond just this particular project, should there be some kind of standing mechanism that if there's some kind of error or perceived error in the root zone contents that that can be reported appropriately. Thanks.

DUANE WESSELS: I agree. Thanks, Kim. So it sounds like to me there's enough changes to warrant a revision to this document, and I guess, to discuss it at a future meeting. Does that sort of match other people's feelings? Yeah. Okay.

So I've taken some notes. I'm happy to have any more questions now, if there are. Or if you want to e-mail myself or the list, I'll take any suggestions there as well. I think that wraps it up.

TIM APRIL: Thanks, Duane. I think we—

DANIEL MIGAULT: I have just one very last comment, which is the introduction of the ZONEMD in that seems to me the least problematic thing compared to the architecture that local root enables, which is coming from a completely managed root zone system to completely unmanaged way to handle the root zone. So I was wondering if that should be somehow mentioned into the document or not at all.

DUANE WESSELS: I guess without some specific text to look at, Daniel, my feeling is that it's out of scope. I think it's an interesting topic. I think it's happening anyway. Maybe it's separate work that RZERC wants to take on, but I don't feel like it belongs in the deployment plan document. Peter?

PETER KOCH: Yeah, I agree with your response. But I would also like to support Daniel's point that the larger issue of potential endorsement or a more in-depth understanding of the consequences of this migration or paradigm shift, so to speak, would be due and probably something that

RZERC should engage in. Again, I agree that that is probably out of scope for this plan.

DUANE WESSELS: Thank you. Tim?

TIM APRIL: Thanks, Duane. I think the only other thing that was on the agenda was to quickly discuss the charter review stuff. Like I said, there's been delays due to external [inaudible]. But I'm hoping to get out an e-mail to everyone in the group of where we are with the discussions about the scoping exercise. I think that we have four topics left that haven't been discussed. But Daniel has been working on setting up the poll for when to have the first meeting for the charter review. So hopefully, we can be in discussion about that and talk about starting that process in the next couple of weeks. Then if everyone can I review the topics that we had discussed over the last couple of months to try and close out whether or not those would be in scope or out of scope, and then we can continue from there. Looking at the last four at some point soon. Any other things to discuss this month? Duane?

DUANE WESSELS: I was going to ask if there are any meetings planned for ICANN73.

DANIELLE RUTHERFORD: There are no meetings planned for ICANN73. Generally, meeting scheduling takes place several months in advance.

DUANE WESSELS: Okay. So it's too late to ask. All right. Thanks.

TIM APRIL: I'm not seeing other hands. Have a good month all. Thanks.

DUANE WESSELS: Thanks. Bye.

[END OF TRANSCRIPTION]