
ICANN67 | Virtual Community Forum – RSSAC Work Session 2
Tuesday, March 10, 2020 – 10:30 to 11:45 CUN

DUANE WESSELS: All right. Hello, everyone. This is Duane. I'll be leading this session on the statement on identification of root server operators. Are we good to go, Steve and Ozan? Do we need to anything before we dive in here?

STEVE SHENG: I think we're ready to go, Duane.

FRED BAKER: Yeah. As far as I know, we're good to go.

DUANE WESSELS: All right. So, welcome, everyone. Thanks for attending this session. The plan today is to discuss this document and a few of the comments that have been raised. The agenda is up there.

I'll go through the rationale for the document and then briefly outline the document itself, talk about some of the comments that we received so far and then open it up for discussion if there's any new comments.

So, this document was originally written by Terry Manderson. I don't believe Terry is on the call but he gets the credit for starting this. He started it I believe back in 2018, quite a while ago, and then it sort of

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.

sat and we didn't do anything with it. Recently, a message went out on the RSSAC that, "Hey, we've got this whole document. If no one wants to take ownership of it, then we'll just abandon it." And I raised my hand to take a leadership role and try to finalize this document. So that's where we're at today.

My recollection is that the need for this document originated, as I said, a couple of years ago. In my mind, when RSSAC started giving these "how it works" tutorial sessions at the ICANN meetings and during the Q&A afterwards, it became apparent that there were these misconceptions out there and that it had become a little bit of a mistake to continue to refer to root server operators by these letters. So, that's where this document came from.

The outline of this document is that ... You can see there's a brief introduction section. The second paragraph describes how it came to be that we referred to root servers as letters, essentially because they were renamed in the mid-90s and it became very common to refer to them as A-root, B-root and so on.

Then the next couple of paragraphs talk about how this led to certain misconceptions over time.

Then the fifth paragraph or so talks about the RSSAC effort, I believe it's RSSAC numbers 0-8 was an investigation into the naming scheme. That didn't recommend any changes, but throughout that work, we did spend a lot of time talking about names for root servers.

Then, lastly, there's a reference to the fact that future architectures may have different ... We may do things differently and need future architectures for the root server system.

And then it ends up with the recommendation to ... In most cases, when referring to a root server operator, just refer to it by name, rather than its letter.

Any discussions or questions about the overall flow of the document before we get into some of the detailed comments?

FRED BAKER:

And let me clarify for my own [value] what you just said about the name of the operator. A.rootservers.net is a name. A-rootservers is a name. By that, you mean instead of a.rootservers.net, just highlighted on the screen in front of us, that would mean Verisign, correct?

DUANE WESSELS:

Yes. Yes, thank you, Fred. That's what I mean. I mean the name of the organization operating the service, not the name of the service. What we're [inaudible] the identity. Not the root server identity. Anything else? Any other questions or clarifications?

Okay. So, up on the screen is the Google Document. This went out on the caucus mailing list a few days ago. We had some comments from Paul Hoffmann, so I've entered Paul's comments here. I think Paul didn't have access to the Google Doc, so I entered them for him. I thought we would just go through those one by one and discuss them.

So, first of all, Paul correctly notes that A-root is not a letter. It's a [inaudible]. Can you click on the "show more" link so we can see his full comment? Okay. So, you can read his full comment there.

But, essentially, he's focused on this sentence where it talks about letters and then it gives as an example the phrase A-root. I think this is something that we can easily fix in editing. It's certainly ... The intention here was to refer to the fact that we use the letter-based names to refer to the root server operators rather than ... So, we call it A-root instead of Verisign, for example. I think this is just an editing thing that we can fix. Anyone disagree or agree on—

FRED BAKER:

I do have a question, and that is that Verisign operates both A and J— do we want to distinguish between those two clusters or not? If we simply say Verisign, we're referring to the entire set of root servers that Verisign operates. And it has two sets.

DUANE WESSELS:

So, I see Paul's hand up. I'll get to you in a second, Paul, since this is your comment. I didn't know you were going to necessarily be on the call. But let me try to address what you're saying, Fred.

The reason that this document exists is because there were people referring to ... When they would say A-root, they really meant Verisign. So, they would say, for example, "How does A-root deploy servers around the globe?" Something like that.

So, certainly, there are cases where when you're referring to the root service, it's appropriate to use the letters. But when referring to the organization, it's not. That's what we're getting at here. If there are times when it's necessary to refer to A-root and J-root collectively, then you could do that if it's appropriate to use those names. But if it's talking about the organization, then you would say Verisign. Does that help, Fred?

FRED BAKER:

Not really. You're making the point that, in the context of the paragraph you're looking at—the paragraph that's right in front of us and has A-root highlighted. You're saying that that's Verisign and that's only half true. If I was to say J-root, that would equivalently be Verisign.

DUANE WESSELS:

I would be happy to take ... This part that's highlighted, these are just examples and I would be happy to take A-root out of the example if it confuses things because Verisign operates two letters. Paul, I think let's get to you since this is your comment.

PAUL HOFFMANN:

Hi. Thank you. So, two comments. First, on what you said earlier about we can clarify this here and you were talking about the second paragraph. I believe the clarification has to happen both here and in the bold face recommendation. That is, when we talk about letters, many people say L, to pick a letter that I like, and many people say L-

root, which we sort of have taken as hyphenated. So, I think that this comment applies both here to the introductory text and should be part of the actual recommendation itself.

DUANE WESSELS: Okay. Of course.

PAUL HOFFMAN: And then the second part is I think I can talk to Fred's question. Fred, I think that ... I mean, yes, we've had this confusion for a very long time, but looking at the very carefully chosen words in this document, which we're trying to make even more carefully chosen. I believe that the recommendation is going to be the thing that is important, and in the recommendation, I think it's fairly clear that there are two contexts for naming a root server. One is the context of the technical context which is where we're now saying instead of using a letter we're going to use the fully qualified domain name, at which point then there isn't an issue between a.rootservers.net and j.rootservers.net. Or when we are talking about operators.

My feeling is ... Again, since I'm only on RSSAC Caucus, not on RSSAC, this is the first I've seen this document. What RSSAC wants is to have those contexts be even more separated, so that in the exact example that you bring up, Fred, which is a root server who happens to have multiple identities, that that's just fine, that that's not a problem, that that's not confusing, so that when we're talking about operators it is correct to use their corporate identities, that that then becomes clear.

I guess, Fred, I would only do this by asking do you want to make it even more clear in this document that we know that today there is at least one entity—Verisign—that operates multiple identities, a.rootservers.net and j.rootservers.net, or do you want to make that even clearer for the future? Because I got the feeling this document is more about future looking.

FRED BAKER: Well, you put that question to me.

PAUL HOFFMANN: I did.

FRED BAKER: I think the context actually is the metrics document—some of the document if it's not the metrics, then I'm losing which one it is—but which refers to the various servers as A-root and B-root and a.rootservers.net and several other things. The first question was could we reduce the number of variations to one thing. And the second question was what should that one thing be?

So, this document came up as a “let's refer to the company” or the organization—it might not be a company. So, I guess at least in part I was just kind of stuck on referring to changing A or J to Verisign. Okay, do we now presume that those are [inaudible] one Verisign service or that Verisign ... Where does that go?

DUANE WESSELS: Certainly not, Fred. You were right that, in the metrics document, the metrics document talks a lot about what we now call RSI (root server identity). So, for example, when you're reporting metrics, you report them as a.rootservers.net, b.rootservers.net, and so on. Nothing in the metrics document would suggest to refer to those as Verisign metrics, for example.

PAUL HOFFMANN: Actually, Duane, I'm going to jump in here because in between the last RSSAC call and this one I actually looked through the latest version of the metrics document, the one that Steve Sheng sent around and that I guess RSSAC just approved.

I'm not sure you're correct about your last statement. I don't think we actually said how the metrics are going to be reported. We said how they are going to be measured, which would certainly be by identity. That is, a.rootservers.net, b.rootservers.net, but we never actually said how they're going to be reported, and even since you were the primary document author there, even your example doesn't say who it's for.

So, having said that, I'm not afraid that we will do the wrong thing. I believe we will do it, we will do the reporting by identity. But I didn't see in the document that it said that the reports were going to be by identity.

DUANE WESSELS: Okay. Well, I was thinking of the example reports as you referenced, and if that got lost, then I lost track of that change. I'm sorry.

PAUL HOFFMANN: I don't think it's that important, but going back to Fred's concern, in anything that comes out of metrics—that kind of metrics collective—I do believe that it will be easy to do it by identity and that people would expect it by identity.

For example—and let's not do Verisign. Let's say that, in the future, ICANN ran two identities, one that's currently called L.rootervers.net and the next one being called X.rootervers.net. It is perfectly believable that ICANN could mess up one of the structures of those two, so that the X.rootervers.net identity actually was doing a bad job with respect to metrics. Therefore, it would be fully expected by the community that those would be split out.

DUANE WESSELS: Yes, I agree.

WES HARDAKER: If I may.

DUANE WESSELS: Go ahead.

WES HARDAKER: Verisign is a root server operator that responds to DNS queries for the root zone at both a.rootervers.net and j.rootervers.net named identifiers. They are still, to follow on Paul now, Verisign is still the root

zone operator and yes it may be operating on two identities and yes it may be having an issue with one of them not serving properly and falling outside of the needed metrics bar. It's still that root server operator that will be shamed—I'm not sure the right word. And fixing that particular identifier service.

UNIDENTIFIED MALE: Right.

WES HARDAKER: So, I think the point of this document is sort of the sentence that I rattled off at the beginning is to do that separation of the organization versus where the organization is serving data.

DUANE WESSELS: Yes. Fred, are you okay if we move on from this?

FRED BAKER: Yeah, go ahead.

DUANE WESSELS: Okay. Certainly, we need to edit this sentence a little bit to fix up this confusion about letter versus letter-root. So, we'll do that. Next comment is, in the same paragraph on the word detracted, it says that these misconceptions detracted from the organizations responsible for providing the service. Paul's comment is that it doesn't say how it

was detracted. And we have another one from Matt agreeing with that. So, we can take that action to provide a little bit more context there.

As I said before, in my mind—I remember this coming up a lot in the Q&A sessions from the tutorials but I can't remember any specific examples at this time because it's been a while.

I would also note that around the time that this document was originally written—2018—we did start to make an effort to be more clear about referring to operators as operators and services via letters, and I think that was successful. In those intervening the years, the times that there have been misconceptions have dropped, in my opinion, so that was successful. But now we're just going through the process of formally documenting it.

DARREN: This is Darren. Do I have [inaudible] to speak?

DUANE WESSELS: Please, go ahead, Darren.

DARREN: I will definitely pass this along to Terry. I'm no longer a caucus member due to inactivity and some health problems last year. I'll probably reapply. But I can pass this along to Terry. But as I recall, some of it was people wanting to have their own letter and the misconception that they could just sign up to host their own letter rather than be a host for one of the different entities that controls the

root server I think is where some of the pressure came from to write this document. But I'll also make sure Terry responds to this.

DUANE WESSELS:

Okay. Thank you very much, Darren. Appreciate that. Any other comments on this sentence and the word “detracted” here? Okay, let's scroll down then.

The next comment, also from Paul, is about future architectures, which he says is a little confusing and proposes referring to not architectures but future changes to RSI naming. I think that's ... I can go either way on this. I would be fine if it says future naming but I also can imagine future architectures that may not fit well into the ... That may be different than the current naming scheme or the current design and would benefit from the separation of organizations and identifiers. Paul, please go ahead.

PAUL HOFFMANN:

So, I'm fine if you want to say future architectures, something that refers to naming. I think that that's fine, too. I don't consider those necessarily architectures. As I said here ... When I think of architecture, like a future architecture, I think of something like let's get rid of the root server system and instead go to HTTPS access and things like that. It seems like every proposal that I think you're thinking about here is actually a naming proposal. So, if you want to leave architectures in, future architectures of the naming of the root servers or whatever, it would be fine with me as well.

DUANE WESSELS: Okay, thanks. Brad?

BRAD VERD: Could you just simplify it even further and just leave it at future changes?

DUANE WESSELS: Sure.

BRAD VERD: I mean, that seems ... I'm a big fan of keep it simple.

DUANE WESSELS: Sure. That all right with you, Paul, if we did that?

PAUL HOFFMANN: Absolutely. I was just typing that in. But yes, for the very reason Brad gave is the fewer words, the less we're restricting ourselves.

DUANE WESSELS: All right. Steve is noting that. Thank you, Steve. When you're ready, you can scroll down. So, the next comment is about the fact that here we have an example, [identity arootservers.net] appears all upper case. It also appears upper case in the introduction. Paul notes that the style that's been adopted is to not upper case these. I think that's perfectly reasonable. I don't see a need for them to remain all upper

case. So, I think that's pretty easily solved, unless anyone wants to argue for why it should remain this way.

Then, the last comment is about what you might consider as missing from this document. There is no table that lists the operators and identities and addresses and so on. This is something that we did discuss, I don't know, a week or so ago. There was a suggestion I believe from Ken to include such a table. We discussed, well, does it make sense to have such a table in this document and also have it, for example, in the history document or should it just be in one place? What happens when there's a change to a root server operator or identity? Do we need to go and update in all the places?

And where that discussion ended up was that we could just reference the IANA URL here which lists the operators. I think part of Paul's comment is that, although the IANA table is good, there are some subtle differences between that table and ... Thanks, Steve, for putting up the IANA page here on the screen.

There are some subtle differences between what appears here and what appears in the draft history document and maybe some other RSSAC documents. So, already, we have the problem of keeping things in synch.

I would like to open up for discussion on whether or not the statement on identification of root server operators should include such a table directly or point to it by reference to somewhere else. Paul?

PAUL HOFFMANN:

I was delaying there hoping to see other hands. Duane, I'm not wanting to be contrary here but I completely disagree with your statement that these are subtle differences. In fact, I would imagine that Wes might be waving his hand wildly here about the subtle difference for b.rootserver.net, where ISI is in parentheses.

I think it's fine to point somewhere else. Again, I'm speaking as a caucus member, not as a root server operator. I think it's fine to point somewhere else, as long as every current root server operator feels that they can change the data at that somewhere else as easily as they need to.

If it was in this document, then yeah, we might need to rev the document with changes later and I think that that's just fine. I happen to be one of the people who thinks that revised documents are fine.

The problem with pointing off to either rootserver.org or the IANA registry is both of those change over time without any change history. So, someone might say, "Oh no, I'm using the right name. Then they go and they realize they're not. So, having it in a document that is revised sometimes I think is better historically because it seems like part of the genesis for this document that we're discussing is that people do care about using their names and I want to make sure that if that is one of the partial genesises that everyone gets to use their name the way they want.

RUSS MUNDY:

So, since you called me out, I totally get your point, Paul. I even agree with it. It's nice that if we want to be the authoritative naming service for the organizations that are participating, it seems like this might be a place that we could still that. I would argue that rootservers.org would be just as good of a place to stick that information.

That being said—and I think that your note on change control is ... Up until that point, I was 100% sure you were wrong. But when you mentioned that history isn't necessarily kept, and we've already shown that in the RSSAC history document how hard it's been to recover past sources of data. That being said, I always get concerned that documents aren't updated rapidly enough and I think that's 1000% true in the ICANN world where we're very slow at getting documents revved. And the root server operators have always, in a previous document, have even said that IANA is the source of the root zone and that is where we get all of our data.

So, I would still prefer, even with a loss of history, not to publish the data in this document. I think the benefits don't outweigh the pain associated with it. And it's absolutely my responsibility to fix the list of root servers on the IANA webpage if I don't like that naming. And you're right. I don't. ISI I think is more prominent than in parentheses. But that's my own fault for not fixing that.

I would pick one of the two places to point to but it sure seems to me if we're making IANA the authority for everything, that's where we should be pointing to.

DUANE WESSELS: Thanks, Russ. Karl, go ahead.

KARL REUSS: Yeah. It was [inaudible] sort of a contrary opinion here. It seems like we spend a lot of time saying what not to call operators in this document, don't refer to letters. But we don't really have any examples of what the operators are. So, maybe not a definitive list, but something along the lines of over time these have been operators of root servers. Just for somebody who is casually browsing the document, I'm not sure they're going to necessarily match the title with what we're trying to say.

DUANE WESSELS: So, Karl, you would like to see the table in this document, it sounds like.

KARL REUSS: Yeah. Although I understand that there's issues with keeping it up to date. So, maybe just referencing some of the operators. But including the table, I'd be for that.

DUANE WESSELS: Okay. So, for those that advocate putting the table in this document, would you also advocate for it remaining in the history document, for example? Then needing to update both places in the future? Paul or Karl, go ahead. Either of you have your hands up, I see.

PAUL HOFFMANN:

I put my hand up here in response to Karl. And also, to be clear, I'm fine if folks want to have a single pointer somewhere else, again, as long as you all understand that that means you have to have change control over there and act on it. I agree with Wes. I think IANA would be the much more logical one.

But for Karl's point, I think that, in fact, we don't need to list examples because I think when you start listing a small set, two out of 13, that actually makes it more confusing. I think we can say not just in a footnote, but in this document, that names have changed, and if you want to see something about that, go look at the history document which we already have a footnote for. I think that that would be a good explanation of why we are not listing names in this document is that names change.

So, what I would want out of this part of the conversation would be something more definitive that says the current set of names can be found at [inaudible]. And in my case, I'm fine with it saying IANA. And say some of the history of name changes can be found in the history document.

DUANE WESSELS:

Okay. So, I feel like we have a little bit of mixed opinions here. I know Karl is advocating for in document, and I think in the chat Darren also. I feel like we don't have agreement on this yet.

UNIDENTIFIED MALE: I was really kind of on the fence. Listening to what Paul just said and reading the last sentence, I'm okay with it as it stands, I think.

DUANE WESSELS: All right. Thanks. So, in the chat, Russ has asked, "Is the table better as a new IANA registry?" I think Naela is on. Naela, are you able to speak about the status of the IANA page that was shown a second ago? Do you consider that a registry or how do you consider it?

NAELA SARAS: Thanks, Duane. Yes, I am here and I can speak. No, we don't really consider this one a registry. By definition of registry in IANA, what's on the IANA webpage is really what we have under the protocol parameters registries. That's what we talk about registries, where there's a process for updating for registries, changing entries in the registries, etc. Per the IETF, the work that we do with the IETF.

This is more on how what we do on the domain name side, an authoritative list of who IANA lists as the authoritative parties that operate these services.

I do acknowledge ... As I was looking at the list, I do see that some of these names ... Yes, I can see some of these are a little bit out of date. For example ... Not to pick on Verisign, but I don't know if Verisign has the capital S in their name anymore or not.

UNIDENTIFIED MALE: Right. It's not preferred.

NAELA SARRAS: Right.

UNIDENTIFIED MALE: Just to be clear, that’s what made me put this comment in initially was when I saw that in the table.

NAELA SARRAS: Absolutely. So, I guess from this discussion, what I can do is maybe take an action item that we’d better communicate the process for changing these names if an operator wishes to update their name. Wes said, “It’s my fault. I didn’t update it.” Wes, perhaps you didn’t know what the process was.

Another thing is I will disagree with you, Paul, what you said earlier in the discussion when you said you have the right to be called whatever you want to be named. That won’t work here, because usually what IANA needs to do is list the name of the legal entity that operates the resource, whether it’s a top-level domain, a root server. It needs to be the legal name of the entity.

So, I don’t know what the legal characteristics are of ISI versus the University of Southern California, so [inaudible] talk about this.

WES HARDAKER: Yeah. That’s a hard one. We are officially the University of Southern California and ISI is a research institute of that. But the University of

Southern California is our legal name. I would have to ask about whether there's an addendum, but that's a side discussion, but thanks for pointing that out.

Paul, thanks for even pointing out this whole problem. I think a lot of us didn't know how badly out of date various things around the Internet were.

NAELA SALLAS:

Yeah. So, I'm in support that there needs to be one sole place for this and it makes sense that it would be the [inaudible] here. I will communicate better about the process to changes, should any operator wish to update their name.

And Russ, if you are asking should this be an IANA registry in the sense of the IETF registries, then I doubt that that's the right structure for it. I think it needs to remain the way it is here.

WES HARDAKER:

Naela, can you answer Paul's other point which is keeping a history of changes with respect to if this table does get modified in the future which I think is likely at this point. Is there any change list that can be published along with it?

NAELA SARRAS:

I will look into that. Typically, we don't have change history available on the IANA.Org website, unfortunately. We've been asked for it in

other places as well. But I can look into that, if we make changes how we make the previous versions available.

WES HARDAKER: That would probably be good. I think Paul’s concern was certainly justified there, although archive.org probably has the changes.

NAELA SARRAS: Right. Yeah. So, is it okay that I take out of this better communication of a process for changes and then change history?

WES HARDAKER: I think that would be excellent.

DUANE WESSELS: Naela, I have one more question for you. Can you go back to that page? The table here uses ... The third column is titled “manager” whereas we would probably call an operator—root server operator. Are you stuck on manager?

NAELA SALLAS: I suspect manager is coming from ... That’s [lexicon] for TLDs. We refer to the registry operators as managers. So let me look into that as well, manager versus ... So, we want to call them operators, correct?

DUANE WESSELS: Well, I don't want to force a certain way of doing things on you but I'm just curious if it's flexible. We could even call it out in the document before us. We could say ... IANA uses the phrase manager while we use the phrase operator.

NAELA SARRAS: Let me look into that one because if it's something that we inherited from other parts IANA.Org where we refer to entities as manager due to the domain names, the TLDs, we maintain ... If that's why we call it manager, then it's something that we can change. So let me look into that before we have to feel like we're stuck with the name manager.

DUANE WESSELS: Okay. Sure.

NAELA SARRAS: Thank you.

DUANE WESSELS: All right. So, I think that's all the comments before us. That's the name of the document. Obviously, there's some work to be done here, some new edits. I'll work on—

BRAD VERD: There's a bunch of hands up.

DUANE WESSELS: I don't see them up. Scroll up. Okay, go ahead. I'm sorry, I was looking at the wrong place. Paul, go ahead, or whoever.

PAUL HOFFMANN: Two things that Naela had just said actually caused me concern, but again I'm not one of the drivers of this document. One is that IANA would only ... Regardless of whether we're pointing to IANA or not, they only put the legal name and that may not be what you all want. Wes just brought up one thing with that. So, that's one concern is if we are pointing at that, does that suffice for what all of the root server operators want, the folks in RSSAC?

I have a very strong issue with a statement that anyone is using saying we're using so-and-so's legal name. The laws of names are very different from country to country and therefore the legal name of Netnod Sweden might be quite different than what the legal name of Netnod is in the United States, looking at Swedish law.

And as a subset of that, unless things have changed in the last 40 years, which they may have, there is no legal name for US Army departments. Don't ask me why I know this but it was part ... At least in 1979, there were no legal names for US Army departments. So, saying that IANA is using that name for that particular root server operator gives me the willies.

DUANE WESSELS: All right. Thanks, Paul. Brad?

BRAD VERD:

A couple of things. One, I'm pretty strongly against having a table in this document. I feel like we should be referencing it.

Going back to the history versus having change history on a webpage, I go back to somebody suggested adding a link to the history document to provide the history there. I'm curious if that suffices that. I feel like we're trying to boil the ocean a bit here.

Then, lastly, I agree really strongly that IANA is a source of truth. The root operators have stated that. We should link to that. We should work with IANA to update this page however we see fit and within their guidelines.

This same page that we're linking to, if you look at the bottom there is a link on this page to the root [hints] file which provides some history. Then the actual names. But of course these names are in a format that we don't like according to the document.

I feel that while there might be some challenges with the IANA webpage, I think those are all fixable, easily addressable I feel. We should give Naela time to go back and get some clarification on that. But I don't believe that we should be creating yet another table that needs to be maintained in a document somewhere. I think that is the wrong directive.

DUANE WESSELS:

All right. Thanks, Brad. Russ?

RUSS MUNDY:

Hi. Yeah, thanks. Responding to Naela, I was not really necessarily thinking of registry and further restricted IETF protocol parameter registries. But in a more general sense of a registry, IANA is actually the registry for the root just as Verisign is the registry for Com and so forth.

So, a general meeting. Having a central point of reference seems like IANA is in fact exactly the right place but the IANA does have certain rules that they have to follow. So, getting a consistent set of both titles and descriptors that are acceptable to both the IANA and the PTI, under their requirements, and the root operators under what the RSSAC sees is appropriate, I think that's probably our broad objective as far as where we point to and that, if we can possibly be consistent between not only the RSSAC documents but between the RSSAC documents and the usage and description of a place that we point to, which is hopefully IANA. Thanks.

DUANE WESSELS:

All right. Thank you, Russ. Any other hands or comments before we wrap up the call? Okay. So, expect a new round of edits to this document. We'll work with Naela to get the answers to those questions that we've asked of her. I'll work with staff to get a new version out the next few weeks.

UNIDENTIFIED MALE:

Duane, when you say get a new version out, will this also come back to the caucus or will this—

DUANE WESSELS: Yes. Yeah, I think it'll definitely go to the caucus. I think it needs to.

UNIDENTIFIED MALE: Thank you.

DUANE WESSELS: All right. Thanks, everyone, for coming. Appreciate your time and I'll see you later.

UNIDENTIFIED MALE: Thanks, Duane.

BRAD VERD: Thanks, Duane.

UNIDENTIFIED MALE: Thanks. Good day, all.

RUSS MUNDY: Thanks, all. Bye-bye.

[END OF TRANSCRIPTION]