/OZAN SAHIN:	Good morning, good afternoon, and good evening. This is the RSSAC 002 Update Call held on the 7th of January 2020 at 14:00 UTC.
	On the call today, we have Dessalegn Yehuala, Fred Baker, Wes Hardaker, Hiro Hotta, Ihtisham Khalid, Jack Biesiadecki, Karl Reuss, Kazunori Fujiwara, Ken Renard, Kevin Wright, Ray Bellis, Roy Arends, Russ Mundy, Shinta Sato, Yoshitaka Aharen.
	From staff, we have Andrew McConachie, Danielle Rutherford; and myself, Ozan Sahin.
	I'd like to remind you all to please state your name before speaking for transcription purposes and thank you. Over to you, Andrew.
ANDREW MCCONACHIE:	Yeah. Thanks, Ozan. So, hopefully, everyone can see the document I'm sharing. This is RSSAC 002v4. There's also been a link posted in chat if you'd rather just go at it through Google Docs. I believe the last time this group met …
	Well, first of all, this is not a formal work party. I'm ICANN staff and this is just an update of a document. So this is kind of a lightweight process that the RSSAC caucus is using to update documents so we don't have to go through all the heavy-handed process of starting a work party and whatnot. So that's why I'm the one talking.
	So, I think the last time we met was November 3rd in Montreal, so it's been a while. The agenda for today is pretty straightforward. I'm just

going to walk through the document and talk about any comments or edits or additional text that has been made since that last meeting on November 3rd.

We didn't completely work through the document at that last meeting. I think we got up through about section 6.7, if memory serves. Another outcome of that meeting, which has not yet resolved, is we have an outstanding action item from Paul Hoffman to investigate YAML versions and I spoke with Paul before this call. I spoke with him yesterday. He's not going to be on this call. It's too early for him. But he is working on that. Otherwise, all the action items from that November 3rd call have been resolved via edits.

So, what we're going to do on this call is just start at the top of the document and work our way through and I'm going to be focusing on things that had been done since November 3rd.

So, the first one is … Section three. Right. This is the first one. Six and three are some measurement parameters. I saw this comment from Russ and I don't quite understand it. I don't know what RSI stands for. Is that maybe something that comes out of the metrics work party or the definitions from RSSAC 26? Russ, can you speak to that?

RUSS MUNDY: I believe it is, Andrew. The term I think we ended up with from the Metrics Work Party, it wasn't instance. It was another I word.

ANDREW MCCONACHIE: Is it identifier?

RUSS MUNDY: I'm sorry?

ANDREW MCCONACHIE: Identifier maybe.

RUSS MUNDY: Yes. identifier. Thank you.

ANDREW MCCONACHIE: Root server identifier measurement parameters. That changes the focus from the organization to the identifier itself.

RUSS MUNDY: Correct. There was still … Although there were several proposed approaches, the one that came to the top was identifier. I think it's most easily thought of as a replacement for saying A-root, B-root, C-root, etc.

ANDREW MCCONACHIE: Okay. Does anyone else have any comments on this before I change it to RSI?

RAY BELLIS: I think it [sounds moot] if we simply don't include it in the first place. This only came up because RSO got added to the header.

ANDREW MCCONACHIE: Yeah. RSO got added to the header because we previously just called it measurement parameters, but then one of the things we decided on November 3rd was to break out the operator—

RAY BELLIS: Yes, okay. [inaudible].

ANDREW MCCONACHIE: Exactly. Yeah. So, we do need a qualifier here.

RAY BELLIS: Okay. Yeah.

ANDREW MCCONACHIE: Are you okay with RSI?

RAY BELLIS: I am, yes.

ANDREW MCCONACHIE: Okay.

WES HARDAKER:	We should certainly put a reference to what that means somewhere earlier in the document, just terminology sake since it's new and unusual.
ANDREW MCCONACHIE:	Would a general reference to RSSAC 026v2 where that term is explained? Maybe we have something like that in the beginning of the document.
WES HARDAKER:	Yeah.
RAY BELLIS:	I personally would rather see an explicit expansion in the document rather than require everybody to go read RSSAC 026.
ANDREW MCCONACHIE:	How about both? I'll do both further out [before] the introduction. Okay. So, let me take an action item to add a reference to RSSAC 026v2 and expand RSI earlier in the document. Okay.
	And as I'm scrolling through the document, if there's something you would like to talk about that I'm skipping over, please speak up.
	The next place I'm going to stop is going to be section 4.1. This is a new comment from Ray I saw you just made today. Do you want to speak a bit more about it?

EN

RAY BELLIS: Not quite sure what else I can say that isn't in the comments. Only important bit is that it's consistent. So, it's really more of a question for the [inaudible] guys as to whether this is a problem for them or not. But yeah, the optimal compressed form of the root zone is not used as [inaudible]. But if they're already reporting using that mechanism, then they should keep on doing so.

ANDREW MCCONACHIE: I don't see Duane on the call.

RAY BELLIS: I see Brad is here.

WES HARDAKER: So, I think the point of that was just to define the standard for how we do it, not how it was transmitted. So, they're independent of how it's actually sent.

RAY BELLIS: Yeah. It was more of a note for implementation or for [influences].

BRAD VERD: This would be a question for Duane and that can be taken offline I think.

RAY BELLIS:	Yeah. I know that, at one point, about two years ago—actually, kind of roughly [inaudible] last updated—[inaudible] changes default. Message size was [inaudible] transfers to be 16k instead of 64k. So, it's something that any implementor just needs to be aware of. But currently, that's [inaudible].
ANDREW MCCONACHIE:	All right. So, I'll take an action item to ask Duane about this and get some feedback from him on it.
	Okay. The next thing is section six at the bottom. Based on discussions from November 3rd, I added this text here, this last bullet. So, I just want to draw people's attention to it because it's new text. It's two sentences down there. If anyone has any comment about that …
	Hearing none, I will move on. Now, there was some discussion about whether or not we should be recording things if their values are zero or if their values do not exist for some reason. I saw, Ray, that you added this text today. Does anyone have any comments on this? This text, which kind of ends the first paragraph of 6.4 is the same in 6.5 and 6.6.
RAY BELLIS:	Yeah. All I've done is actually rephrased it, so I don't think the original where it says only size ranges with non-zero counts must be listed. It's kind of an [inaudible] conversion. So, I've literally just rephrased that to say they must be omitted.

ANDREW MCCONACHIE: Right. That's more clear, actually.

RAY BELLIS: Yeah. It's just wordsmithing. There's no functional change.

ANDREW MCCONACHIE: I'd be tempting to use "should" but that's a nit. My IETF hat coming in.

RUSS MUNDY: Well, yeah, but I actually understand, just reading them, and I haven't thought about this before but they seem to be opposites of each other to say that only list things that have non-zero counts. Well, okay, yeah. So, size ranges with a zero count must be omitted. Okay. Okay. I retract.

RAY BELLIS: Yes, correct. That's exactly why I was rephrasing it, because I think the original phrasing was unclear.

RUSS MUNDY: [So be it].

ANDREW MCCONACHIE: So, I'm going to move down and here we have the same thing with 6.5. Same change. Then …

RAY BELLIS: Yeah, 6.6 should probably have it as well.

ANDREW MCCONACHIE: Yeah. [inaudible] sources. Did I spell omitted right? Yes, okay. Yes.

RUSS MUNDY: But I think he meant … Yeah, okay.

ANDREW MCCONACHIE: Yeah, that's better. Thank you. Okay. So, this is as far as we got November 3rd. We didn't get to 6.7. I think we stopped somewhere in section six, slightly before this. So, there's a new proposal that came up on the RSSAC Caucus list to lengthen the URL path standard. So, currently the URL path standard, it can kind of start anywhere in the directory tree.

So, my basic understanding of URLs is you have the domain name and then after that, you have a directory tree. And the proposal that came up on the caucus list was to expand the standard all the way back to the domain name part of the URL, so that you would have domain name and then you would immediately start with what is specified here as opposed to having some kind of maybe some leading directory immediately after the domain name.

I haven't made any edits to the document reflecting that yet because I wanted to see what the caucus had to say about it before I made any changes to it.

RAY BELLIS: I disagree with that proposal. I think every—and this is with my RSO hat on as opposed to a caucus member hat on. Every RSO publishes their base URL which can include a path component on the rootservers.org website. There is absolutely no need to have a consistent initial prefix of the path, in my own opinion. And I note that we do not have … Sorry. We, F-root, do currently have a [inaudible] prefix of RSSAC 002 within all of our paths.

WES HARDAKER: I agree with Ray. The ending path needs to be consistent but the prefix does not. In fact, such a specification wouldn't even get past in the IETF without [inaudible] or something.

ANDREW MCCONACHIE: Does anyone have any further comments? I'll just make no changes here. The one change that has been made here is the RZM has been specified for the root zone manager for the [short service]. I assume that's kind of non-controversial. I don't think there was anything specified before for that.

RAY BELLIS: That's not controversial. Is there somewhere prior to this point where [short service] is actually defined for everybody else? I haven't noticed it yet.

ANDREW MCCONACHIE: [inaudible] short service is a shorter version of the service name. Generally, [inaudible].

RAY BELLIS: So, it's only specified as an example currently, not formally. Okay, sorry. It's [inaudible]. Sorry. [inaudible] letter-dash-root. Okay, sorry. It's right there. I couldn't see it [inaudible].

ANDREW MCCONACHIE: No problem.

RUSS MUNDY: So, question. Did anybody respond on the caucus list to [inaudible], the guy that made that comment about the path [format]?

ANDREW MCCONACHIE: Yoshitaka made the comment about how we should specify RZM short service. I don't remember who made the comment about lengthening the path.

RUSS MUNDY: Okay. But my fundamental question is did anyone respond to him or her?

ANDREW MCCONACHIE: I don't know. I would have to go back and look at emails, unless someone remembers.

RAY BELLIS: Yoshitaka is on the call.

RUSS MUNDY: Okay. Well, Andrew, since you're running the show here, could you go back and look at the email and find out who made the path? [inaudible] comment and make sure that they got a reply. It just seems inappropriate to someone made a comment and have it not go into the document and have it not sorted out on the list.

ANDREW MCCONACHIE: Yeah, that's true. We should give them a response. I will do that.

RUSS MUNDY: Okay.

ANDREW MCCONACHIE: Okay. So, the next edit here is from me. This is change to recommendation three that was suggested on the caucus list. Previously it said in two years, and based on discussions on the caucus list, I changed that to "as necessary" so that it doesn't specify a specific timeframe in which the RSSAC must update this document. Any commentary on that?

RUSS MUNDY: Well, speaking for myself, I agree with it.

RAY BELLIS: Yeah. It's okay. My own concern is that [inaudible] does it specify a [inaudible]. I think the intention was to say, yet, we don't want to say we must [inaudible] nothing has happened, but conversely, [as written], this is now saying we should do something if something big happens in six months. I don't know if that's the intent. If everyone is happy with that, that's fine.

ANDREW MCCONACHIE: I think the intent was to avoid having to update this document when not much had changed and having to do a bunch of work when there was no real need to do the work. That's how I interpreted the intent, at least.

I mean, we can certainly change this language to say something like should be revisited every three years or as necessary or something like that. I mean, we could qualify it. We could put a minimum, so it wouldn't be revisited until every three years if that would work.

RUSS MUNDY: Andrew, would it make sense to say that it should be reviewed every X number of years since the last revision to determine if updates are required? That doesn't require an update if it doesn't need it.

BRAD VERD: Well, you can just leave it how it was then, right? I mean, the intent of this document or the intent of that change was to not require us or

EN

require us to go through all these documents and update them if they weren't necessary, but we could put it however the group wants to do it. We have a lot of these documents that kind of have this in the end of them and a lot of them don't change.

RUSS MUNDY: Yeah. I think one of the concerns from some folks, Brad—and I guess I'm kind of in that camp—is with the "as necessary" statement it could just be basically ignored for a long time without something that said at least look at it to see if it needs an update.

BRAD VERD: I have no objection to it. This was brought up as an issue so we were trying to address the issue. I'm just channeling the issues as they were brought up.

ANDREW MCCONACHIE: Let me try some wordsmithing right here.

RUSS MUNDY: I think right there I would use the word "whether" rather than "if".

ANDREW MCCONACHIE: How do people feel about that new text?

RUSS MUNDY: That looks good to me.

RAY BELLIS: Looks good to me.

ANDREW MCCONACHIE: Okay, so we'll leave that as is. Okay, another new comment. I just saw your comment today on this, Ray. We can certainly include the caucus members who helped put this document together, the previous versions [inaudible]. I don't imagine anyone would have an objection to that, but if you do have an objection to that, please say so now.

So then under acknowledges, the RSSAC Caucus members will list the previous members who helped put this document together as well as the new members who helped put this document together. They'll all just be listed in alphabetical order.

I had to go through and renumber things, so that's why these are all updated. The only new text left is this. This will likely change. Ken did a good job of putting this together, the version four update. I believe it's accurate but this will probably again require a bit of revision when the document is finally done, so I wouldn't take this section as gospel at this time.

Other than that, I am out of things to talk about. Does anyone else have any comments about this document? I think my plan would be to merge everything we've talked about and to make the edits that we've talked about on this call and to talk to Paul again about investigating which YAML version works best and maybe add some text about that into the document. But otherwise I don't think there are any outstanding issues

with this document. Does anyone have any issue with this document they'd like to raise at this time that we haven't discussed on this call?

Okay, this is going to be a short call, guys. I guess I can give you 30 minutes of your life back. Thank you very much for participating. I'm going to make these edits and talk to Paul and I would expect the new version … I think I can send a new version to the caucus mailing list probably maybe later this week or early next week. Thanks, everyone, for your input. Go ahead, Brad.

BRAD VERD: I was just going to say thank you for your work on this Andrew.

UNIDENTIFIED MALE: Thank you, Andrew.

ANDREW MCCONACHIE: You're very welcome.

RUSS MUNDY: Thank you.

JACK BEISIADECKI: Hello. I had a quick question about clarification. Say in the event of … I'll use load time as an example. If there is no data for that, should that date be posted or should it just be omitted completely?

EN

ANDREW MCCONACHIE: Good question. Let me go down to that section of the document. Yeah, load time. So, here it says, if we're talking about load time, 6.1, if no load time metric is available to be represented as [the integer -1].

JACK BEISIADECKI: So, for example, that 2016010100: instead of 811 it should be -1?

ANDREW MCCONACHIE: Correct. That's how I understand it.

RAY BELLIS: In practice, I'm not sure how easy that is to implement. To the first update of the day, that's fine. You can [inaudible] there's been one every day. But what happens if you don't see the first one or the second one? How do you know [inaudible] if you never saw it?

ANDREW MCCONACHIE: That's a good question. It's something that I've come across and that's why I was asking the question, to figure out how to proceed in a scenario where, say, hypothetically there is no data for a particular day.

RAY BELLIS: Yeah. I think you're right to [put] this up, actually, and that more thought needs to be given because technically we are supposed to be reporting serials and we shouldn't assume a serial ever actually existed until it does exist. I don't know how we prove that.

ANDREW MCCONACHIE: Do we need more discussion in this section on this topic? It sounds like just putting a -1 there is a bit naïve. It's not something we can actually implement in all instances, so maybe we need a bit more discussion there.

RAY BELLIS: Yeah. I'm trying to think how [inaudible] actually works because it's so long since I wrote it, but I look at the logs in each of our servers and I look at the serial numbers that are there and I find the first instance [inaudible] and I look at when that [inaudible] on the other servers and then calculate the time differences and then [inaudible] percentile.

But if, for whatever reason, a serial number never appears, I'm never going to report it. We'll see when that happens. That's actually ostensibly an error in zone distribution but that's something to be picked up by the metrics work, vantage points more so than for this. So, I'm inclined to basically say that you can't report anything [inaudible] serial number you've not seen.

ANDREW MCCONACHIE: Yeah. Would it make sense then to just not list that serial number as opposed to listing a -1 next to it? If it's a serial number you haven't seen, then you just don't include it.

RAY BELLIS: I think so. I mean, by all means, create the file, but it would have no timestamps in it.

ANDREW MCCONACHIE: Yeah.

RAY BELLIS: Because we shouldn't be assigning semantic value to the time stamps for the purposes of this. I mean, yes, it happens. Verisign always used Y, Y, Y and then [D, D, N, N]. Apart from that, as far as we're concerned, they're just [inaudible] labels.

ANDREW MCCONACHIE: We can change the final sentence to just say if no load time metrics is available, it should not be listed. And you may end up with a file with no lines in it, basically, no time entries.

RAY BELLIS: Yeah. That's feasible. Actually, looking at this now, that doesn't look like valid YAML in that example there. Sorry, just [inaudible]. I think the two dates should be indented. But yes, you could end up with a file that has no entries in the time array.

ANDREW MCCONACHIE: Hopefully, I fixed the YAML and I'm wondering if you're okay with that sentence as well.

RAY BELLIS: The double [inaudible] is a little bit [inaudible] convoluted.

ANDREW MCCONACHIE: No load time metric is available, okay.

RAY BELLIS: It's lightly [inaudible]. So, what's the [inaudible] three, if anything?

ANDREW MCCONACHIE: I'll take a look.

RAY BELLIS: Yeah. I'm looking for it now. It says if no load time metric is available, it should be marked with dash.

ANDREW MCCONACHIE: Well, that's just as problematic as marking it with negative one.

RAY BELLIS: It is, yes. Ultimately, we can only include those serial numbers we've actually seen appear on systems.

ANDREW MCCONACHIE: Well, let's think about this a bit more and maybe take a suggestion to the list. For right now, I'll leave the text as this. If you get a chance to

look at your implementations and seeing what you're actually doing there—

RAY BELLIS: I know that my implementation will basically not include any serial number that hasn't [been observed]. I'm pretty sure that's the case, which is probably not [inaudible] complaint with what [inaudible] said. But I should double check that.

ANDREW MCCONACHIE: Well, it may be compliant with version four.

RAY BELLIS: Well, yes.

ANDREW MCCONACHIE: Okay. Well, let's leave this as an open item then for now. I don't want to take a bunch of time on the call wordsmithing and whatnot here. Thank you. Thanks, Jack, for bringing that up. That's a really good point. Yeah, thank you.

Before I close out the call, does anyone have anything else they'd like to bring up? Any issues with the document?

Okay. With that, I'm going to end the call then. Thank you, everyone, for all your contributions. You'll see a mail from me on the list later this week or next. Thanks.

UNIDENTIFIED MALE: Thanks, Andrew.

RUSS MUNDY: Thanks, Andrew.

[END OF TRANSCRIPTION]