

OZAN SAHIN: Good morning, good afternoon, and good evening, this is the RSS Metrics Work Party teleconference held on the 17th of December 2019 at 1700 UTC. On the call today we have Duane Wessels, Ken Renard, Abdulkarim Oloyede, Fred Baker, Jeff Osborn, Kazunori Fujiwara, Kevin Wright, Ray Bellis, Russ Mundy, Ryan Stephenson, Shinta Sato, and Dessalegn Yehuala. From Staff we have Steve Sheng, Andrew McConachie and myself, Ozan Sahin. I would like to remind you all to please state your names before speaking for transcription purposes. Thank you, and over back to you, Duane.

DUANE WESSELS: Alright, thank you very much, Ozan. This is Duane. We have just a few things to go through today regarding the metrics document. To keep you up to date a little bit about what's happened recently, a couple weeks ago we made some final changes to the document, we locked the document for edits and sent a PDF copy around on the mailing list for any feedback on those last changes. We received a few comments on the list and then yesterday Russ and myself met with Steve to go through those comments and do a short editing pass, so we did that.

Hopefully, this call today will be the last work party call. Then we will do an RSSAC caucus last call with a deadline of January 10th. And if all goes well, then the RSSAC could vote for approval as early as their February meeting. So, that's the tentative plan and I'll walk you through some of the most recent changes and a couple places where new text has been added, just to point that out. So, that's the plan.

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.

Ozan, can you take us to the end of the introduction. I think everyone has had a chance to see this. This last paragraph in the introduction is something that was added, it's probably been about a month or so, now. But I just want to make sure that people have had a chance to look at this.

Like I said, this has been in the document for a while so hopefully everyone has had a chance to read this, but this is kind of an important addition. So, if you haven't seen that, please take a look at that last paragraph in the introduction. Next, Ozan, let's go to Section 3.3. I need to pay attention to the chat. Did you have a comment, Ray, or somebody?

RAY BELLIS:

Sorry, I did, but I was trying not to talk too much as I've got a bad cold. I was just rereading some stuff, and I think in context this new paragraph, in the very first sentence of the introduction, where it says this defines metrics to ensure RSOs are making minimal performance, I don't think that text is actually correct anymore.

DUANE WESSELS:

Okay. Alright, let me make a note of that and we'll address that. Alright, thank you Ray. So, let's go back to Section 3.3 then, Ozan, if you are there. So, Section 3.3, we're talking about vantage points and this section is about connectivity and other requirements. This middle paragraph has been sort of highlighted by Paul as possibly being unnecessary, and we did get some feedback about this on the list. Our

sense was that most people were in favor of keeping this paragraph in or keeping something very much like it in the document.

I should also explain perhaps why I feel this is important, why this was added. It's because when Paul was doing his proof of concept implementation of the metrics he had access to VMs all from a single VM provider and so he put his code on 8 or 10, or so, virtual machines around the network and his data showed that a couple of the root server operators had extremely good latency on the order of 1 or 2 msec.

So, I had a concern that the choice of vantage point connectivity provider could lead to that kind of situation where it just so happens that a particular operator or identity has very good connectivity to all of the vantage points. So, this paragraph was added to ensure some amount of diversity in connectivity providers.

I think Paul's point here is that when you choose a vantage point you may not actually know which connectivity provider is handling traffic. You may go to something like an exchange point where there is already lots of connectivity providers and it may not be possible to specify which ISP the packets flow out of. But I think for the most part comments were in support of keeping this, and that's kind of what Russ and myself are advocating for. But I'd like to open it up for discussion if anyone has any thoughts at this time on this paragraph.

Alright, not seeing any hands, and by the way, if you're not able to raise a hand, just feel free to speak up if you like. Okay, so not hearing discussion on this, let's move on to the next one. Ozan, let's go to

Section 4.9. This is where we talk about the K/n system model and there is this formula. Ray, you had earlier left a comment here that this formula or this approach doesn't really take into account the way at least certain resolver implementations behave, and again, I think we had some discussions on the list about this.

My opinion was while that may be true, we're sort of adopting a very simple model here and we didn't really want to get too much into the way that different implementations work. I think the others supported that, as well. And so with sort of your permission, Ray, we would like to leave this as is. Any comments on this section at this time from the work party? Okay, again, I'm not seeing any hands so we'll move on. Let's go to Section 5.3.

So there are a couple new sentences here, and this is again related to something that Ray brought up in a previous call where a name server instance when it first starts up it might have been turned off for a while or disconnected for a while, so when it first starts up it may have a stale zone file and if it happens that a measurement occurs between the time where it starts up with the stale zone file and it's able to refresh the zone data, then the measurement may get stale data.

This particularly could impact the correctness metric because the correctness has a rule that data must be within 48 hours and due to the way the root zone is signed, the name server can start up with data that's as old as seven days and still be allowed to serve it per the parameters of the SOA record. So this new text here acknowledges that that's a possibility but leaves it for future improvement or leaves it to see if it actually becomes a problem in practice.

There is nothing here that proposes trying to work around this fact of the way the name server software works. So again, Ray, since this is yours, if you have any thoughts on this addition, I'd love to hear them.
[CROSSTALK]

The next references the exact same thing, so this possibility also affects the publication metric as well. So, here's the text in Section 5.4, that just sort of references that earlier discussion.

RAY BELLIS: This new section of 5.4 doesn't quite stand right.

DUANE WESSELS: It needs some editing work, you're saying?

RAY BELLIS: Yes, it's the bit before current zone has been refreshed.

DUANE WESSELS: Okay, yes, thank you. Alright, so we'll take a look at that. Any other comments about this? Okay. Then let's go to, a little bit farther down, I don't think there's a marker here in the text anymore, Ozan, but again, this is the publication latency metric and there was a suggestion at some point to maybe change from using median here to use the mean of the aggregated measurements.

So, the thought was that there were some maybe somewhat contrived examples where a root server identity could have slightly under half of

its publication latency measurements exceed the threshold, yet still meet the metric and switching to mean would solve that. And again, I think we had some comments about this on the list, maybe in the document, as well. And I sense that most people were happy or satisfied with leaving it as median at least for now. They felt that the example in which you had these sort of extreme values was sort of, as I said, contrived and unlikely to happen in practice.

My comment was two things, one, I'm a little bit opposed to using mean because this would be the only place that we had a mean aggregation in the document and I'd like to keep things consistent, but also there are other ways to address the contrived example or the concern of having outliers. Instead of median we could use some other percentile value.

So, after our staff call yesterday with Russ and Steve, we proposed to leave this as median, as it is here. I had a private conversation with Paul Hoffman who said he was okay with median, as well. I think he had originally raised this. So, unless somebody strongly wants to advocate for mean at this time, we'll keep it as is. So, now's your chance, any comments on mean versus median here?

Okay, I'm not seeing any hands or discussion, so I think we're good on that. That's really all the things that are sort of outstanding or pending in the document at this time. As I said, the plan is to do a last call, give the caucus a deadline of January 10th to make last comments. So we expect people to give this another read-through perhaps, maybe find things that we have overlooked, things that don't make sense due to all the changes, or even grammatical errors, expect to see those, and then we'll go back to the RSSAC for the final vote.

Steve, Ozan, or Staff, anything that I missed? Or Russ, any last points to make before we wrap up? Fred, go ahead.

FRED BAKER: Are you proposing a call on or about the 10th?

DUANE WESSELS: We are not proposing a call. What we'll do is we'll just do it over email. Steve will send out a PDF and ask for comments to come back in email. If there are significant problems found that need discussion, then I guess we would have another call, but assuming there is not, there would not be another work party call. Does that sound okay?

FRED BAKER: Yeah, that's fine, I just wondered what you were thinking.

DUANE WESSELS: Okay .

RUSS MUNDY: I just came up with what is potentially a modification to that first sentence in the introduction. I just put it into the chat room.

DUANE WESSELS: Oh, this is where it said ensure, is that what you're talking about, Russ?

RUSS MUNDY: Where it said ensure, right, and talked about minimal level and so forth. So it's sort of a quick modification, but if we could take a look at it and see if it addresses Ray's concerns and would work for the rest of the group, we could do that substitution.

DUANE WESSELS: That's a good idea. Why don't you read it out in case not everyone has access to the chat.

RUSS MUNDY: Sure, so the current intro paragraph sentence in the first bullet underneath Introduction says, "defines measurements and metrics to ensure root server operators are meeting a minimum level of performance." And the proposal that I just came up with, "defines measurements and metrics that root server operators meet to provide an appropriate level of performance." I wanted to keep inference out of that, especially in view of the paragraph what was added with respect to the contracts and so forth.

DUANE WESSELS: Russ, I like your change, I like that it changes from passive to active, or past tense to present tense, I think that's good. Regarding must or should, I think you're right. Again, since this is just sort of the introduction, I think avoiding must and should is probably a good idea.

RUSS MUNDY: So, if folks are okay with that change, we'll put that in before we send it out for the last call, then? Ken, are you okay with that? Okay, good. So, people, relook at it, think a little more about it, please comment on it in the last call, too.

DUANE WESSELS: Yeah, we'll take care of that. I guess, Russ, you and I and Steve should talk about whether or not, maybe we'll send out two versions, like a red line version and a clean version for people to comment on, something like that.

RUSS MUNDY: Yeah, that's probably good.

DUANE WESSELS: Okay, yeah, thanks for doing that on the spot, Russ. Anything else, folks? Okay, well, thanks everyone for making it today and thanks for having a short meeting we us. I see some more text in the chat. Naveed says we might need to define appropriate. Yeah, I missed that you had changed minimum to appropriate in your sentence, may we leave it as minimum?

RUSS MUNDY: Well, I think Ray did have a comment on that later on, but I'm thinking possibly if appropriate isn't the right word, perhaps I wouldn't have a problem with minimum, but I think that Ray did identify that that

doesn't quite fit the context of the document at this point and what I'm think, again on the fly here, is "expected."

SPEAKER:

That would work, to a large extent, though, appropriate or expected are all subjective, but the metrics are defining what that is, so appropriate is defined by the result of the document.

DUANE WESSELS:

Yeah, right, that's kind of what I'm struggling with, too. Again, since this is the introduction, I don't think we need to get into a lot of detail about what is appropriate, at least not at this point, that comes later. I don't know. I'm on the fence. I can see leaving it, just say the root server operators meet a minimum level of performance, but I'm fine with appropriate, as well.

Well, I feel like we shouldn't maybe wordsmith it right now on the call. So, let's put the change in the document and we'll leave a comment about if folks have a preference or if there are concerns about appropriate or minimum, and have the discussion in the last call phase or on the list before then. Okay, well, with that I think we'll wrap up the call and everyone have a great holiday season and we'll see you after the New Year.

[END OF TRANSCRIPTION]