
ICANN72 | Virtual Annual General Meeting – RSSAC Work Session
Thursday, October 28, 2021 – 09:00 to 10:00 PDT

OZAN SAHIN: Hello, and welcome to the RSSAC Work Session on updating RSSAC047. My name is Ozan Sahin and I'm the remote participation manager for this session.

Please note that this session is being recorded and follows the ICANN Expected Standards of Behavior. During this session, questions or comments submitted in the chat will only be read aloud if put in the proper form as noted in the chat. I will read questions and comments aloud during the time set by the chair or moderator of the session.

If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor. Please state your name for the record and speak clearly at a reasonable pace. Mute your microphone when you're done speaking.

This session includes automated real-time transcription. Please note this transcript is not official or authoritative. To read the real-time transcription, click on the closed caption button in the Zoom toolbar. With that, I will hand the floor over to the RSSAC047v2 Work Party leader, Anupam Agrawal.

ANUPAM AGRAWAL: Thank you, Ozan, for the context setting. Then we move into the roll call.

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 SAHIN: Anupam, I think at this time, we will take the attendance from Zoom Room. I think we can proceed to agenda item two.

ANUPAM AGRAWAL: Thank you, Ozan. We move to agenda item two, which is review of the decisions/action items from last meeting. Steve, can I request you to take this agenda item forward?

STEVE SHENG: Thank you, Anupam. I'm happy to do that. Good morning, good day, everyone. Our last work party meeting was on 13th of October at 16:00 UTC. We started going over the list of issues that Paul has put together for RSSAC047v2.

One of the decisions and action items that Warren has, as part of his implementation, identify some issues and he will add this to the list of issues as a running list for the document. I don't see Warren on the call. I reached out to him early in the day. So I'll work with him to add those if he hasn't already done.

When the work party go into the discussion regarding the issues related to elapsed times and timeouts, the original discussion is whether to account for TCP setup time. But a discussion went into how to correctly account for the latency of flipping over UDP to TCP. And the work party had a discussion, I think the high level decision points were that the purpose of the latency measurement is really to

establish the round trip time. It is not to measure truncation. It's not to measure it from the end user perspective. So, nevertheless, these points will be added to version 3 discussion.

The other major discussion point, the work party engaged was the procedure for handling outstanding issues that are out of scope for version 2 but potentially for future versions. And the decision there is really, first of all, this work party should focus on a narrowly scoped version 2 statement of work, which is aimed to correct any technical errors identified in the document and to put a timeline for every few years to review the recommendations in the document. But nevertheless, the work party agreed to create a new section on the current issues list. That will list as the work party discuss issues arise in version 2. Some of those are out of scope of version 2 that might be useful to be tackled in a future version. So to the extent the work party will list off and at the end of going through the current issues, the work party will have a discussion of these parking lot items and potentially to scope out version 3 a bit and give it back to RSSAC for their consideration.

So that's kind of the decisions and action items from the last teleconference. Any questions? If not, let me give it back to Anupam. Thanks.

ANUPAM AGRAWAL:

We have Paul's hand. Yes, Paul, go ahead.

PAUL HOFFMAN:

Thank you, Anupam. I spoke with Warren at length earlier this week or late last week—I forget—privately about his issue list. So even though he’s not here on this call, I expect that he will participate. And at this point, as far as I understand, all of his issues would be in the parking lot. We don’t have any specific issues that would go in the v2 document, although as he works on his own implementation—just to be clear, I mean, I hate speaking for somebody else, but his implementation is not a direct implementation of the requirements on 047. It’s taking many of the ideas that were in 047, and doing the implementation, and then sort of shoehorning them towards the actual requirements. So I’m hoping that he does find some things in that shoehorning process that are not clear in 047 that we can use as clarifications. But at this point, I’m comfortable with us moving forward with the current list of v2 things, knowing that some might get added before we’re done and that many of what Warren has. And I have a bunch of issues as well for the parking lot that we can do those at the end instead of trying to run them too much in parallel. Thank you.

ANUPAM AGRAWAL:

Thank you, Paul. I think that’s the right way forward, considering the scope, which has been given for this work party. So thank you for the comment. We move to item three now.

WES HARDAKER: A quick interruption, if I may, I've managed to get a hold of Warren. He will join this call in a few minutes. It'll take him a minute. So you may want to go on other things. He'll get here.

ANUPAM AGRAWAL: I see him now. There he is.

WARREN KUMARI: Hello. What did I do this time? Am I in trouble?

ANUPAM AGRAWAL: No, you're not. We were just talking about the point where we discussed last time that some of the issues which you have identified as part of your implementation, you are going to add it in to the running list. And we were having a comment from Paul saying that you and Paul had some discussion and most of your issues are going to be part of the parking lot, which is going to be a new section within the document for the scoping of version 3. So do you want to add anything to that?

WARREN KUMARI: Not really other than—I mean, I still have some concerns or disquiet of how we account for certain types of behaviors. Like if we're reporting the latency at a very high level of precision then we need to include things like what the actual or at least my belief is we need to account for things like when queries fall back to requiring a second query because of behavior like truncation, whether that should be recorded

as part of the latency. Let me restart that because that was no way coherent.

So we're recording things like latency reporting with them at relatively high precision. To me, that doesn't seem as though it makes a huge amount of sense if we don't do things like account, we don't include the full transaction. So for example, if F-root is answering all of their queries in 50 milliseconds and X-root answers them in 30 milliseconds, but every query requires two round trips because you always get a truncation, then it feels as though the metric is completely meaningless, and we shouldn't report with that level of precision. We should just say it came in under four seconds or not. Basically, if we report the metric, it should have actual meaning and not be completely synthetic. But I have a bunch of other comments as well. But those are more for long discussions during the working group chat.

ANUPAM AGRAWAL: Yes, Paul? Go ahead.

PAUL HOFFMAN: So, Warren, one of the things that was said just before you touched in was that my proposal is that we actually wait on parking lot issues until we've gone through the purely technical issues. Maybe a quick note on them. One thing that you bring up, though, that's interesting, because from our discussion the last time we decided that that would be a significant technical change, and therefore, not related really to

the v2 document. But one thing that you bring up which is interesting is we sort of made the queries that are currently in RSSAC047 be unlikely to be truncated. That is, the two types of queries are the .soa queries and certain pre-formatted queries for correctness. All of those are unlikely to be truncated even at 512 bytes.

I think it would be interesting to note, even if we aren't going to change the timing, is that the metrics collector could note whether the TC bit was set in the responses. So that would be a way for us to figure out in the future how to deal with this. I believe with our current set of queries, they would all be—no [inaudible]. But if I'm wrong, that is actually quite relevant for future discussions.

WARREN KUMARI:

I believe they are not all without TC bit because some set of operators appear to be doing the—it's completely skipped out of my head—but if you get the same set of queries from a resolver, you bucket them and then cause them—it's an [inaudible] thing. Paul came up with it. Brain completely shut down.

PAUL HOFFMAN:

Okay. Thank you. Again, I think that's interesting for us to look for on that. And certainly all responses, we would be able to tell if there's a TC bit. So even though that's not part of RSSAC047 currently, I think that's a reasonable addition to add of look for this and see what that is. I certainly can do that in my example setup. I suppose you can as well. So thank you.

ANUPAM AGRAWAL: Ray, you have your hand up.

RAY BELLIS: Hi. Warren, the techniques you were mentioning is responsible for RRL?

WARREN KUMARI: Data. Thank you.

RAY BELLIS: I don't believe—sorry. Someone is echoing badly. Possibly yourself. Can you mute for a second, please, Warren?

WARREN KUMARI: Yes. Just trying to find—

RAY BELLIS: That's much better. Thank you. So there's RRL. And I don't think that the queries that we'll be seeing as a result of this measurement will be anywhere near fast enough to trigger RRL. We're normally looking at multiple queries in space per second to even start triggering that. That's going to be the case. I would say that if we do decide to measure, if you'd like, the so-called end user experience by measuring second trip, we must also at least record the first trip as well and not just the total. I wouldn't give this off without a risk to find a route leak

from very fine grade external monitoring which dropped—yeah, a very simple routing configuration change drops average latency in Europe from 22 milliseconds to 15 milliseconds. And that was kind of the basis of this sort of external monitoring. We wouldn't have otherwise suspected this was happening. So if we've been measuring the total round trip time, we would never have seen that. By looking at the very fine RTT measurements, that's actually something that we saw and we fixed and we're better for it. And that's actually why in the early versions document, we actually said it's important that operators be able to see the figures that they're getting and because those might actually have an actionable intelligence for them.

ANUPAM AGRAWAL: Okay, Warren. You're on mute, Warren.

WARREN KUMARI: [Inaudible] I'm responding to Paul. An example is the responses of QNAME equals TLD and type equals ns. That not infrequently seems to exceed 512 bytes. Not drilled down, see whether that's the main cause of the TCs that I see. Currently my audio is choppy. But yeah, I will try and do that instead.

ANUPAM AGRAWAL: Okay. Fair enough. Ozan, can we move to the top of the document to see? So we had some amount of discussion on this elapsed time and timeout. Any comments on this? Warren, go ahead. You're on mute.

WARREN KUMARI: Sorry—

PAUL HOFFMAN: I think that we can resolve this by saying that even though in my first initial implementation, we were using dig and people seemed to think dig was okay. This limitation of dig, in fact, is not inherent. And as I think I said at the last meeting, I was starting to use, for example, DNS Python as a replacement, where in DNS Python for the TCP, you can actually set up a socket and then do the request over the socket. So you can time the socket set up and the whole setup.

So I think that we might be able to finish this one by saying, “Yes, you really need to use a tool that, in fact, does the measurements the way we want.” I think that it would be a perfectly reasonable conclusion here because we now know or at least I now know, certainly other people knew before me, that such tools are available and can be used.

ANUPAM AGRAWAL: Thank you. Any other comment on this one? Do we agree with this? Yes, Duane?

DUANE WESSELS: Follow up to Paul. Paul, it sounds like there’s really no change needed to the document for this particular issue at this time, right?

PAUL HOFFMAN: Correct. So no change to the document, just that if in the future somebody says, “Oh, this is difficult,” the answer is “Well, okay. Change tools.”

DUANE WESSELS: Yeah, okay.

ANUPAM AGRAWAL: Warren?

WARREN KUMARI: Quick question. Hopefully my audio works. Also it doesn’t really say what sort of resolution we should be reporting it. The way I’m doing it is I do my measurements and I use Prometheus, which bucket stuff into I think 10 milliseconds, [inaudible] like 15 or 20 buckets. And visually, that seems to show a relatively good distribution. If we’re only monitoring or if we’re only bothering to report with a precision of 100 milliseconds then a lot of these problems become simpler. Now, if we’re dealing with microsecond precision, then the measurement methodology needs to change because you need to decide, do you measure it when the packet leaves your Ethernet interface or what? Obviously micro standard, we would be crazy pants territory. But there is no real discussion here on what precision to record at.

ANUPAM AGRAWAL: I’m having a little difficulty in understanding because there is some noise in the audio. So yes, Paul, if you can move ahead.

PAUL HOFFMAN: Yes, I've spoken with Warren enough over the past to understand him even through a lot of static, so he can put on his cute headset. I understood what Warren said and I just did a quick check of the document. He's completely correct. Somewhere in Section 4, we should probably add a statement that says the precision is within milliseconds. But we don't actually say milliseconds starting when exactly an ending when. But I believe if we say milliseconds, the measurements are all done in milliseconds with an assumed error of no more than 10 milliseconds for startup and shutdown, that that would be good. So this would be a new addition to Section 4.

ANUPAM AGRAWAL: I think that's first addition also. Warren, is this okay if we had this couple of links in the document?

WARREN KUMARI: Let me make sure audio is working, one sec.

ANUPAM AGRAWAL: Much better.

WARREN KUMARI: Okay, good. Yeah, I think so. Actually, Paul had said—we mentioned milliseconds but it doesn't actually say with millisecond resolution, as far as I see it. An example, 27 milliseconds response latency. It doesn't

actually say that all is done with that. But yeah, having what the actual time that one's measuring is from when it leaves the OS, when it leaves the tool, etc. Sorry, I think I just rambled. Whatever you said, it's fine.

PAUL HOFFMAN: But, Warren, I have a question for you on that, which is, since we will be measuring from an application and the application may have a certain startup time for the query that might be within or not, is that reasonable to say with an error margin of 10 milliseconds?

WARREN KUMARI: I would think so.

PAUL HOFFMAN: Okay, because what I was seeing in my current implementation where I am using DNS Python, instead of dig, the TCP startup time even to quite distant places seems to be on the order of magnitude of two or three milliseconds to start a TCP session, which means it's going to get swamped by the error. But I don't know that we can pretend to be any better than that.

WARREN KUMARI: If people don't mind, I could quickly show a graph, which I think would help explain why I'm asking.

PAUL HOFFMAN: I think that that would be difficult during this. And also, when you showed me your graphs, they were hard to see.

WARREN KUMARI: Okay.

PAUL HOFFMAN: But please do send screenshots to the mailing list. Because once you described it to me, I thought that was in fact interesting.

WARREN KUMARI: I will send pictures.

PAUL HOFFMAN: In the chat, Kazunori actually makes a very good statement that the drill command, which many people use, includes the TCP setup time. So I think without us trying to get too precise, we still can say approximately what precision we are going for. And I would say with a 10 millisecond error would be from the sort of back of the envelope looking that I'm doing would be good. But maybe we can discuss this on the mailing list more.

ANUPAM AGRAWAL: Paul, Ray writes in the chat that "If a distant TCP handshake is completing in sub 3ms, it's not being measured correctly."

PAUL HOFFMAN: Okay. That's very interesting. So yes, let's take this to the mailing list or the question of, "Is the 10 milliseconds a reasonable starting thing?" And once I get some better numbers—I, unfortunately, keep changing my setup so I start over again constantly—I can report those.

ANUPAM AGRAWAL: That would be good. Okay. Thanks, Steve, for updating it. Can we move to the next item?

STEVE SHENG: Thanks, Paul. We'd like to discuss this, the TCP Fast Open. Thanks.

PAUL HOFFMAN: Sure. If you scroll down a bit where we have the note, that would be good. In short, the 047 says must not use TCP Fast Open but is not specific enough because there's two TCP—anything in TCP normally is done in the kernel. And then for TCP Fast Open, an application needs to ask for it. So it's not clear if the requirement here is make sure that TCP Fast Open is turned off in the kernel or make sure that the application that is doing the measurements is not asking for TCP Fast Open.

So I would say that we need to add a little bit of wording saying that. And either is fine. That is, in all kernels, it is pretty trivial to either turn off TCP Fast Open or turn it on. By the way, current Linux kernels, for example, that I'm using in the vantage points have TCP Fast Open turned on by default. But both of the tools that I used, dig earlier and

DNS Python now, do not turn it on and all. In the case of dig, you would have to ask for it to be turned on in a certain way, and in DNS Python, it's not even currently available.

ANUPAM AGRAWAL: So what is the suggested change, Paul? You want to include this information specifically now?

PAUL HOFFMAN: I would like to hear from other people about, given the current wording, how do we want to improve the wording to be clear where we are saying—where in the implementation is TCP Fast Open to be prohibited? Is it in the operating system or in the application or in both? Right now it doesn't say and I don't know whether we want to go into greater detail. But I want to hear from other folks on that.

WES HARDAKER: Why doesn't either work, Paul?

PAUL HOFFMAN: I think either is in fact just fine. Wes, are you saying that you would change the and/or in the second sentence there?

WES HARDAKER: I am not a TCP expert. That is not my lair.

PAUL HOFFMAN: Me neither. Absolutely. I had to ask a transport person about this.

WES HARDAKER: But it seems to me that if TCP Fast Open as a problem, it doesn't matter where it's turned off as long as the net effect is the same.

PAUL HOFFMAN: Okay.

WARREN KUMARI: What we care about, I believe, is we want TCP Fast Open to not be used for the query. How that's actually implemented doesn't matter, right?

PAUL HOFFMAN: Okay.

ANUPAM AGRAWAL: Okay, Duane, you go next.

DUANE WESSELS: Wes and Warren said what I would say, which is it should be either. I don't particularly have issues with the current language. But if Paul or someone has suggestions for different language, then we should consider it. But in my opinion, it's fine as stated.

ANUPAM AGRAWAL: Okay. I saw Fred’s hand briefly. Fred, would you like to comment here?
Or is there a comment taken care of?

FRED BAKER: Well, I was in the process of writing it in the chat instead. I think it
might be sufficient if we’re going to measure and if we’re going to
mention this, to mention where we want to measure it. If we see a TCP
Fast Open on the wire then it has been used. And before that point,
does it really matter? You might mention that. I’m done.

ANUPAM AGRAWAL: Okay.

PAUL HOFFMAN: I think that where I’m having an issue is that—I don’t understand the
third sentence in this. So I understand the second one, which is none
of these features should be turned on, great. But environments or
operating systems that do not allow TFO to be disabled should not be
used for the measurements, if at all possible. That’s not necessary if in
fact the application is not asking for TCP Fast Open.

So, for example, Warren’s doing his implementation. He may be using
an operating system that he just doesn’t know how to turn off TCP
Fast Opening, but it doesn’t matter if you he’s not asking for it in the
application setting the queries. Possibly removing that last sentence
and simply leaving that none of those features should be turned on,

and we can even change that saying these features must not be turned on on the measurement platform. Does that sound right?

WES HARDAKER:

Let me throw out another alternative, which is the reason I raised my hand a bit ago, which is we should really have an understanding of the level of Fast Opens that are used to in like an OARC sample or in a DITL sample so that we can figure out what it is we're supposed to measure. If we're trying to measure what are we trying to come up with stats that are used to measure what clients actually need, if we have a 50/50 split, then we should be measuring with and without Fast Open. If nobody ever uses Fast Open, then we should be measuring without, right? So maybe measuring both might be the necessary change to this document, if it turns out that there's a wide mix of usage.

PAUL HOFFMAN:

Wes, that to me sounds like a parking lot issue, though, because the current document says none of these features should be turned on. And I hear what you're saying and I'm not disagreeing with it for parking lot for v3 of, well, maybe we should have it turned on sometimes so that we are getting interesting measurements. But I don't think that we want to be doing that in the current document, given that RSSAC in the work party statement said only technical clarifications and such. Whereas that I think what you said—and I'm just double checking—would be a technical change, yes?

WES HARDAKER: If the old document said, “Don’t use Fast Open,” then yes, that would be a technical change.

PAUL HOFFMAN: Okay.

WES HARDAKER: I don’t remember what the document said.

PAUL HOFFMAN: If you look on the screen, where the three sentences are there starting with some features such as.

ANUPAM AGRAWAL: This is the old paragraph.

PAUL HOFFMAN: Yeah. So the way that this document, the Google Doc that we’re editing is, is it has the actual wording from 047, followed by my notes of what should we do about that? So I think let’s push Wes’s idea for the parking lot later. And I think it is an interesting one, not just about TCP Fast Open but about any modern changes that we are concerned with is that before we do v3, in between v2 and v3, we may want to do some additional measurements to see whether there are new relevant things to do. And some of that certainly could also go back to what

Warren had said in the earlier part of if we're seeing a bunch of TC bits set, for example, how do we want to deal with that as well? So I think that's a more generic parking lot issue but we already have two items for it, which is TC bit being set and TFO. Thanks.

ANUPAM AGRAWAL: So, considering the discussion at the part of the paragraph which is if at all possible. So there was a comment from Fred that this is a very open-ended statement. I think the intent of the first paragraph was that none of these features should be turned on and environments or operating system that do not allow TFO to be disabled should not be used. Now, this "if at all possible," is this something which we can take it out, remove it?

WARREN KUMARI: Yes. I do not know of any implementation where you cannot force off TFO, and if it is, we just won't use that implementation.

ANUPAM AGRAWAL: Okay. Any other comments? Yes, Duane? Do you want to add something more?

DUANE WESSELS: Well, I wanted to respond to Wes and push back a little bit on Wes's suggestion. What I heard Wes's saying was something like the nature of the queries coming from the measurement system should sort of be similar to real world traffic and to sort of measure what the user

experience is. And I want to push back on that. I don't think that the purpose of RSSAC047 is to measure the user experience. It's to ensure that the root server operators are meeting certain service level expectations. In this document we've sort of said, well, UDP and TCP have different characteristics and we want to measure those differently. TCP Fast Open sort of complicates this because it makes TCP latency look a little bit more like UDP latency. And the only way that we can get consistent and comparable results is to disable TCP Fast Open and other things that reduce connection setup delays.

WES HARDAKER:

To respond, I agree, Duane. The job of the document is to measure the components of the RSS, absolutely and the service that is provided by those components. My thinking was that those measurements should tie to what is expected and maybe our notion of what is expected should shift if we find that we are measuring something that is not used. Why are we trying to establish a service level minimum for a TCP option that is or isn't used? When the reality is, I don't know, I just want to measure the right thing or provide an expectation of a level of service based on some real world data so that we know what's important.

DUANE WESSELS:

Yeah, I agree that it's always good to base decisions and policies on real world data. My feeling is—and again, this is just a feeling. I don't have data to back it up. But my feeling is that TCP Fast Open is not

prevalent enough today to warrant using it in this measurement platform, but maybe we need data to back that up.

WES HARDAKER: I was just looking into seeing if I could get data quickly. And of course, it's not quick to do.

WARREN KUMARI: Apologies for the potentially broken audio again. But I think some of it also comes down to continuing our discussion on the mission resolution, TCP Fast Open as a somewhat minor thing. If we have an error bounds then that it's less important than if we're trying to get exactly that.

ANUPAM AGRAWAL: Yes, Paul, go ahead. We will bank on your experience to understand Warren.

PAUL HOFFMAN: I believe what Warren just said was that the reason why we specified don't use Fast Open in 047 was it would speed up some things, but it might be speeding them up by so little that they're within the error bounds, which means measuring or not is not that important.

I raised my hand, though, to say this is all parking lot material. This is not about 047 at this point. I think it's valuable parking lot material. And like I said earlier, I think it's interesting that we are now—not that

we're filling up the parking lot. That I totally knew we would do, but that we are finding things that go in the parking lot that are related. And so Warren's statement that the reason why we cared about TCP Fast Open is because of latency and such and that that will matter I think is good. But again, I think that this is a parking lot item. Once we get to the parking lot in a future meeting, I think it will be valuable for us to look at which parking lot items are related to each other and try to group them, especially if they're going to be something where we can't decide what we want in v3 just by looking the parking lot items. We might need to do some research before then. And so an interesting 2.5 would be test beds that are looking at parking lot items for v3.

ANUPAM AGRAWAL:

Makes sense. Thank you. Any other comments on this before we move to the next item? Ozan, can we move a little ahead? Yes. So this is 4.5 on spoofing protections.

PAUL HOFFMAN:

So this section was on taking reasonable steps to preventing acceptance of spoofed responses and how to do that. One of the things in this list, which I think was put there by Ray so I'm hoping Ray puts his hand up fairly soon, was using the mixed case. My note here was that—I easily could be remembering wrong but during the workshop that we had before finishing 047, we discussed mixed case and decided not to use it because we wanted RSOs to be able to, in their logs, search for the queries from the vantage points. And yet, we've left it in here. So, this is a note saying that I have not done mixed

case randomization in my implementation, even though it is specified here. If, in fact, everyone agrees, it should be there. That's quite trivial for me to do. But I didn't do that because I remembered people said the vantage points want to be able to look where they are being probed.

ANUPAM AGRAWAL: Yes, Ray, your comment, please.

RAY BELLIS: I have no recollection of me being one that suggested that.

PAUL HOFFMAN: Sorry for accusing you of that. Do you have an opinion, nonetheless?

RAY BELLIS: That one especially? In terms of searching for the queries, it doesn't matter to me whether they're mixed case or not. Pretty much what they're talking of cases to search.

PAUL HOFFMAN: Okay. Well, instead of pointing at Ray since I was incorrect, let me ask, do anyone here have an opinion on searchability of logs versus using mixed case as a spoofing protection? Because we know that mix case is a reasonable spoofing protection. And it is pretty trivial to do in any vantage point collector software.

ANUPAM AGRAWAL: Yes, Warren. Go ahead.

WARREN KUMARI: A whole bunch of our queries are [inaudible] like dot or a TLD which is relatively short. So our queries are not really unique enough that people can realistically search for them by string anyway, right? People need to be able to search by IP address, [inaudible] or dot queries or SOA dot, etc. So I suspect that the mixed case randomization is not going to be particularly useful either way. I don't think it's particularly useful for spoofing protection because it's not particularly long set of strings we set up [inaudible] as they say are just draft. But also, because if people would want to find out where [inaudible] then they don't any at all that they would need to search by IP so I don't think this is one that we specify either way.

PAUL HOFFMAN: Okay. To do Warren interpretation, Warren says that he didn't feel strongly either way on this because he didn't feel like it's good spoofing protection, but he also didn't feel like it was that searching by string would be that valuable. I think this was mostly there for the 10% correctness cases. As Warren pointed out, those are the only ones—everything else is sort of for TLDs anyways.

So I would say let's take this to the list as something to see whether anyone really does not want, still does not, or whether my memory is correct, that people did not want the mixed case randomization. And if

not, then we can take out this note and simply say, yes, it is still one of the spoofing protections needed. I see Duane put up his hand though, so please join.

DUANE WESSELS: I don't remember any discussion about mixed case in the context of searching for things. So since it's optional here, I think it's fine. I think your implementation can do it or not, your choice as the implementer. I don't think—

PAUL HOFFMAN: It's actually not optional. Or the way I read this where it says vantage point software must use proper source. That sentence starts with a must.

DUANE WESSELS: Yeah, but it says optional. Optional "Ox20" mixed case.

PAUL HOFFMAN: Oh, I see. Okay. Yes.

DUANE WESSELS: So in my opinion, no change needed to the document. You can do whatever you want.

PAUL HOFFMAN: Very good. Thanks.

ANUPAM AGRAWAL: Thank you, Duane. Any other comments on this? Seeing none, can we move to 4.8? Yes, Paul. What is your recommendation here?

PAUL HOFFMAN: I believe we won't get here in 12 minutes. There is a lot of can clauses, as I point out here, and dealing with unexpected results. In my implementation, I have found a bunch of unexpected results. I bug Duane with them on a regular basis. Now I can start bugging Warren as well. And I don't know really how to deal with this in the sense that the things in these four paragraphs really would need to be done as a collective of the RSOs, which is a future thing. It would require RSOs to come to consensus on after seeing this, we come to this decision.

So I don't have anything in my implementation currently that handles the ability for an RSO to actually question things. I mean, it can be done informally. And maybe the answer is we just forget about this for now, that this is only really valuable in the final implementation. But I was a bit concerned because, in fact, once my initial implementation is working better, which it will be, or, for example, if Warren finishes his initial implementation, or anyone in the work party, I encourage other people to do their own implementations or at least poke at mine. Start doing that. And data is getting published to the RSOs. What do we do when an RSO—the purpose of 4.8 is for RSOs to be able to question what's going on when it looks like something is wrong. What do we do about an RSO saying, "I believe something's going wrong"? This to me feels like it might actually be part not of 047 but of

the future group that will be determining things. So I'm not sure. Like I said, I flagged this as in I did not put any of this in my implementation and I don't know if anyone expects that. But if they do, I would need a lot more guidance.

ANUPAM AGRAWAL: Yes, Duane. Please come in.

DUANE WESSELS: Thanks, Paul. I guess I wouldn't really expect your initial implementation to really have anything to do here. As you said, I think a lot of this would be for a future root server governance, body, or committee to sort of deal with on a case-by-case basis. I think for these situations, there would need to be humans involved in reviewing specific situations and making decisions on excluding measurements when appropriate, but I don't think there's anything for you to do in your implementation and if that's helpful.

PAUL HOFFMAN: I think that is. I assume you mean just not my own implementations because now we have Warren and, again, hopefully others. In any initial implementation, you would not expect these paragraphs to have some way to be implemented. You don't feel like there would be a need for a way for an RSO who's looking at the initial data saying, "I object" or "This is unexpected" or things like that. Is that true?

DUANE WESSELS: I think that's true. Yeah. It would be interesting if we can get to the point where your implementation and Warren's implementation are regularly producing data, maybe even producing the sort of reports that are envisioned by RSSAC047, and then roots server operators can say, "Oh hey, look, on this month, my report wasn't good. I failed. Let's go look and understand what happened." And we can sort of exercise that whole process like you have to go back and look at the actual data and realize, "Oh, maybe this vantage point was offline at this time." But I would consider that all just sort of practice for the real root server governance bodies.

PAUL HOFFMAN: I would then maybe as a thought for the parking lot is that for v3, we would actually remove this section because it's not about—it is expressing a concern but it is not actually about how the data is collected or processed. It's more what will happen after that. So that's just a thought is that for the parking lot is maybe remove section 4.8 from the document. Not to say that it's not legitimate but it's not appropriate in a document that is talking about the metrics themselves.

DUANE WESSELS: Well, I guess we can have that discussion later. But I don't necessarily agree with that. But that's a later discussion.

PAUL HOFFMAN: Okay.

ANUPAM AGRAWAL: Yes. Steve, you can flag it as for the discussion in the next meeting. Warren, you have your hand up.

WARREN KUMARI: Thank you. I can see two things. One, I had thought that if the report says that a root completely missed all of their metrics for an entire month anyway, we could stand and be like, “Whoa, no.” It turns out that this doesn’t make sense to be [inaudible] and they didn’t bother to do any of this stuff. So I thought this was more of auto reports are generated or at report generation time, it’s like [inaudible] the fact that we missed the metrics because I thought that was more of a manual thing. Apologies [inaudible] each time as my example.

For my implementation, I think I’ve mentioned that most of what I do using a letter which generates stats and shoves them into Prometheus. I had also implemented a bit where when I see incorrect results or anomalies, I stick them in Elasticsearch. One way that we could do some of this is I could take a random sampling of the results and always stuff them into Elasticsearch or similar with a tag to say it’s a random subset. And that way, if in the future, the root server is to say, “I disagree with your metrics on Thursday at 3:00PM. I could say, “Here’s a random sampling showing the actual query that was sent, the result that was received, the latency and timing, and why I believe the metric is not whatever you believe it is.” If the tool were to actually use something like Elasticsearch or something like that, presumably the operators themselves could also have access to that and look at

the data themselves. That's actually kind of related to one of my other questions of we're building a system. It's unclear if the output is just supposed to be, "Here's a report from the system," and it says, "X-root had this number at the time," or if we're going to the trouble of build it all, the operators should have access to look at the metrics in real time and see them. Maybe it's specified somewhere in the document but my assumption was like, if I'm collecting it and stuffing it in a system or the [inaudible] access to look at the metrics probably at least for latency and general status across the entire system. And then for their particular instance, much more detailed stats like here's a collection of queries that I thought are incorrect or here's a collection of latencies for your server from different vantage points or something like that. Maybe that's [inaudible].

ANUPAM AGRAWAL:

Okay. Warren, thank you. Yes, Paul?

PAUL HOFFMAN:

So just briefly, because I know we're about to run out of time. This was again a thing that was discussed during the development of 047, which is the purpose of this, to do the monthly reports, or is the purpose of this to also let root server operators see the incremental data so that they can fix things. This is definitely a parking lot discussion again. The reason why I pulled this out here was to point out that there are different ways—the document 047 is completely clear on how to make monthly reports. It is not clear on how to do anything in the interim or the motivation or how that should go and

such like that. So, really, my purpose here was to say if there's going to be interim data or even during a month, whereas Warren said, a root fails everything, then there needs to be, I believe, either a lot more or a lot less text here. And my proposal was less, Duane disagreed, which I think is fine, but again, this is a parking lot discussion.

STEVE SHENG:

Thank you, Paul. I think we are running to the top of the hour, but I'd like to reserve some time for Ozan to talk about scheduling next steps before handing it back to Anupam. Ozan?

OZAN SAHIN:

Thank you, Steve. Hello, everyone. Although we noted that the next call would be on Wednesday, 10th of November 2021, in fact, I have not sent out the calendar invitation for this call yet. So my question is, since we don't have any other calendar invitations for upcoming calls, shall I send the calendar invitations for call every other week again on Wednesday until the end of the year? That would give four more work party calls to this group, I guess. And if this is more than needed, we can always cancel the extra calls. So that's my first point.

The second one is the work party has met again every other Thursday at 16:00 UTC so far. I note that some countries are coming to the end of the daylight savings time. So as we move forward, do we want to keep the 16:00 UTC time slot, or do we want to shift to the 17:00 UTC on Wednesdays? Thank you.

ANUPAM AGRAWAL: I see Warren. Do you have something on these timelines?

WARREN KUMARI: No. I had my hand up from before. I guess I just note that next week is not the week after that because IETF, I believe, at least some of the people involved in this will be at many of those.

ANUPAM AGRAWAL: Ozan, we have a suggestion in the chat for 17th November.

OZAN SAHIN: Yes. Thank you, Paul. It makes sense to skip the IETF week. We can start the calendar invitations afterwards. What about the time slot? Do we want to keep 16:00 UTC or do we want to go to 17:00 UTC on Wednesdays?

ANUPAM AGRAWAL: 16:00 appears to be fine. So if there are no other objections to 16:00, we can continue with 16:00.

OZAN SAHIN: Thank you, Anupam. We'll do.

ANUPAM AGRAWAL: Okay. Thank you, everyone, for being on today's call. We will see you again on 17th of November. Thank you. Bye-bye and have a good ICANN 72, rest of the meetings.

STEVE SHENG: Thank you, Anupam. Bye all.

OZAN SAHIN: Tech support, please stop the recording.

[END OF TRANSCRIPTION]