ICANN70 | Virtual Community Forum - RSSAC Work Session 3
Thursday, March 25, 2021 – 09:00 to 10:00 EST

OZAN SAHIN: Hello, and welcome to the Root Server System Adviser Committee, RSSAC work session three. My name is Ozan Sahin and I'm the remote participation manager for the 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 view the real-time transcription, click on the closed-caption button in the Zoom toolbar. With that, I will hand the floor over to the Work Party leader Abdulkarim Oloyede. Abdulkarim? Abdulkarim, you're muted if you're speaking. I still cannot hear Abdulkarim at the moment. Ken, would you like to take the floor?

KEN RENARD: Sure. Thank you. If Abdulkarim does come back, I'd like to invite him to say some opening remarks. He's the Work Party leader here. So, welcome to the Work Party meeting for the local perspective tool. Since

we have a wider audience than the normal work session for this Work Party, I wanted to go a little bit of the background and history of this Work Party, catch people up to speed and then we'll dive into some current work on the document. So, the Work Party started about one year ago.

ABDULKARIM OLOYEDE: Can you hear me now?

KEN RENARD: There's Abdulkarim.

ABDULKARIM OLOYEDE: Sorry about that. I lost connection briefly there.

KEN RENARD: Yeah. So, Abdulkarim, if you want to make some opening remarks, then I could go into some background because we have a wider audience today.

ABDULKARIM OLOYEDE: Okay, that's fine. Thank you. Let me, first of all, welcome everybody to the local perspective working party monthly meeting which is holding during the ICANN70 online meeting. This working party has so far had a monthly meeting since the inaugural meeting which was held in April 2020. The aim of this working party is to develop tools and metrics which are needed in measuring the local perspective in order to help

improve the Internet services in general and the root server system in particular. And this is necessary because [inaudible] points on the Internet where the service level is not as good as others for one reason or the other. Therefore, this working party was set up so that the RSSAC can have a tool or like a set of tools that can easily measure the local perspective on the root server system at various location points of the Internet.

The tool would help in measuring and identifying the inadequate infrastructure and identifying locations for the new root server system instance and this will also help give the end-user a measurement of their perspective of the RSS. During our early meetings, we have different presentations and originally, we developed seven narratives which were reduced to five and this was also finally reduced to three use cases. And the three use cases we have now is number one is information, the determination of underserved areas, evaluating the third party request to host an Anycast instance and also, we are looking at the recursive server operators.

These narratives would help us to identify the possible scenarios, thereby identifying the relevant metrics or measurements that is needed. And the working party is winding down at the moment. We're winding down our activities and we're finalizing on the document to be delivered to the RSSAC. I want to use this opportunity to thank everyone who has contributed one way or the other to this document. More especially, I have to thank the core writing group which has usually been meeting outside in normal meeting of these working party. I want to especially thank Ken for his leadership role in this Work

Party and I also want to thank RSSAC for giving me the opportunity to lead this Work Party. I think I would now hand over to Ken to go through the documents. Thank you very much.

KEN RENARD: Thank you, Abdulkarim. Expanding a little bit on what Abdulkarim described here, so this tool—this local perspective is somewhat of a complement to RSSAC047 which was a Work Party to measure the performance of the root server system and RSOs, root server operators themselves. So, whereas that work focuses on performance of the RSOs and are they fulfilling their requirements and kind of health of the root server system, really measuring from widely connected areas and getting that perspective, this compliments that work where we're now moving that measurement out to the leaf areas of the Internet where a site, an organization, a small network enterprise can see, well, what does the root server system look like from here?

As Abdulkarim went through the user narratives, informing the decision of an underserved area, so we're not defining what underserved means. That's a very subjective concept. We're trying to collect information that can be used in the decision of determining whether an area is underserved. The other use case, the second use case for RSOs receiving a third-party request to host, this tool is meant to assist those third parties and root server operators to gather information to see what the perspective looks like from there, whether it's worthwhile for an RSO to place an instance at that site. This is not a requirement that an RSO use this tool. It's just hopefully a tool that they could use if there

are additional things an RSO wants to collect technically—technical measurements, they could be added to this tool. It's more of a way to share our measurements and information that can then feed a very subjective decision process.

The third scenario where recursive operators really just share some of the perspective, some of the other user narratives that we combined into this one, if you're anywhere on the Internet and you wanted to see what does the RSS look like from here? Is it good? Is it bad? These are measurements that can inform that decision. So, for each of these user narratives, we've discussed what the user might be interested in, what measurements are of interest and important. Previously, we had some very detailed ways to analyze the data. So, take this measurement divided by that measurement, take this, lots of fancy math and stuff like that.

We got rid of that. We were just being too much engineering here and less document and wider perspective things. So, our use case here, our goal is to determine what information is necessary and how to measure that to inform these decision processes. In the document, we then go on to summarize the measurements from each narrative roughly into pseudocode so that we can almost outline exactly what a tool would do. So, we've done that by expressing the measurements in terms of common tools like Dig or traceroute or ping. So, not that those tools are specifically required, just that it's a common syntax that we can express exactly what we need to do.

These measurements are not particularly unique or groundbreaking. Most of these measurements overlap with what Atlas probes already do and we think that that is a good thing. If all the measurements we're looking for are already included in Atlas probes, then we have instantly about 10,000 measurement points. That's fantastic. So, however we can, we want to leverage all the good work done by RIPE in the Atlas Project. But for a user, a point on the Internet where you want to get this perspective, get these measurements to inform other decisions, this tool could be easier or more lightweight than deploying an Atlas probe to that area. So, the use case could be, again, working with an RSO or just somebody that just wants to see, is curious and wants to do their own analysis of the data. So, we are describing a tool and what it needs to do. We are not necessarily implementing it. Initially, we had started off with the great goals of implementing this for every platform on earth, mobile phones and everything, but that quickly became out of scope.

So, as participants here today, we invite you to contribute to the document, to the project, either by bringing up ideas in this meeting, if you're a caucus member, you're welcome to participate via the mail list, the caucus mail list or directly into the Google document here today. But again, note that this Work Party is wrapping up. We're not taking large conceptual changes to it. Anything that you do want to contribute is welcome but we ask that you do the writeup and we'll see as a group whether it's worth including at this late stage of the game. With that, I'd like to go into current work, going through the document and continuing with the Work Party effort on that. Are there any questions

or comments on the background? Not seeing any hands. Ozan or Andrew, if you're monitoring the Q&A, I'm looking at the chat too but if somebody could bring that up to me? Fred comments about the title of the document. Yes. please, we're open to suggestions here on the actual title of the document. Yeah, Fred?

FRED BAKER: Well, okay. So, let me say verbally what I'm getting at in the chat. If I look at the title page of the document, Andrew has a question basically saying, "Someone, please suggest a better title." And the thing is that I think the title is actually backwards. It talks about measuring the local perspective of the root server system of tells me that attributes the perspective to the RSS. I should think that it's a user's perspective for a—somebody's perspective, looking at the root server system. So, this is literally a question, would that be a better title?

KEN RENARD: I don't particularly have a preference but this is fine. What we're really measuring is the local perspective.

FRED BAKER: Well, yeah, if I'm a company or a country or whatever and I'm using the root server system, the question is which server is going to be most useful to me and/or which set of three servers.

KEN RENARD: Well, I think we want to—which set of three servers or which servers are most interesting to you is a way of getting at your view of the entire RSS.

FRED BAKER: Well, yeah, if I'm looking at 13 different identities, which is the one I'm going to use right now? And I'm looking for the information to make that decision.

ANDREW MCCONACHIE: I changed your measurement to plural. Otherwise I think it's a fine title. I mean, I'll just go ahead and accept it. I don't hear anyone disagreeing with it.

KEN RENARD: Sounds good. Thanks, Fred. Ray?

RAY BELLIS: Shouldn’t it be the local perspective rather than a local perspective?

ANDREW MCCONACHIE: Are there multiple local perspectives?

FRED BAKER: Each person looking at the root server system has a perspective on it. That's why I said, a, as opposed to the.

RAY BELLIS: It's the implicit third person in this case with respect to current user.

ANDREW MCCONACHIE: Yeah. I actually don't know which one is right. I mean, we can go with the, maybe I hear what you're saying, Ray. From the perspective of a user using this tool, there is only one local perspective.

RAY BELLIS: Yes. That's a better way what I was trying to say.

FRED BAKER: Okay. So, if it's going to be the, there, I think the word “of” needs to change to “on.” It's, you're looking at the root server system. No, no. Yeah, there you go. That particular word.

KEN RENARD: Okay. As we go through this document further, again, we want to, yeah, the document is becoming more stable and we are wrapping up this Work Party in the next month or two so, again, changes like that on the title are absolutely welcome and encouraged. But if you've got something more major of a change, you will need to present text for consideration. So, if we move down, I guess, Andrew, to section 2.2, I had an editorial rewrite of this section header trying to align it more with the rest of the flow of the document. Ray and I chatted and he had a comment. I don't know if you can see my comments. I don't know why my comments—I have to press the button to submit it, don't I? So, if you can expand the comment that I just put in there.

FRED BAKER: So, I have a meta question. In front of the document, I see a microphone symbol saying that the sound is up. Where is that? I don't think that's on my computer. I think that's being displayed.

KEN RENARD: Where is that?

WES HARDAKER: It's your computer, Fred.

FRED BAKER: It's my computer? Okay. Well ...

KEN RENARD: Okay. For section 2.2, I can try to do an editorial rewrite of this. One of the things I just did recently, there's a sentence in here. It hopefully captures the concept that Ray was trying to get across. So, saying, well, each RSO will have different criteria on whether to accept such a proposal, including many non-technical factors. It is desired to have a set of measurements. So, if Ray, you could take a look at this, if this captures what we're saying, others as well take a look at this comment to essentially replace the current three paragraphs that it—

RAY BELLIS: Yeah, it does capture that right here. Given the book mentions and non-technical factors, yeah, that's fine.

KEN RENARD: Okay. I will hold off accepting that change or making that change, inviting others to spend the time to read it and comment on it. Okay. The next part was in the core writing group, we had talked about potentially removing the entire third user narrative on the recursive operator so that's why it's stricken out here. So, this user narrative was maybe some organization, if they had a, you know, measurements of the RSS using this tool from several points on their network, that could potentially inform them of maybe routing changes that they would want to make or maybe even architectural changes if they can get better root server system performance by moving a recursive server here or there. Now, the note that previously we had a user narrative on like a government or organizational perspective. I think the use case Abdulkarim had brought this up of an area that wanted to see how well this was performing via their network. So, that was sort of rolled into this recursive operator narrative as well so I just remembered that since the last meeting and wanted to see what people think about this. The other two narratives are really informing RSOs. So, this was the only one left to really give the user community or the Internet community in a broad sense a narrative, at least in this document to describe what they could see of the RSS. So, Fred and Ray, you have your hands up and if they're old hands or new hands, Fred, go first. Okay, Ray?

RAY BELLIS: Yeah, hi. As soon as I was the [inaudible] objected to this in the last meeting we had in the core working party, actually, I think the way you

just reframed that with one exception, I think would actually make this more palatable in terms of, for example, looking at things like routing changes they might want to make. The bit that always got me with this particular writeup at the moment was when [it talks] about the placements for the new recursive server. I think frankly, that's [bunk]. I don't think anybody would ever change the position of a recursive server or deploy a new recursive server just because they have this tool. But they would make routing changes potentially. And that's something we see all the time with our own peering that, you know, we see something suboptimal, we change routing but we don't change placements. So, I think we can reframe this in terms of the IP routing layer and at least so we've got some visibility of where things are and get rid of this idea about placing a new recursive server. I'd be quite happy with this. But that's also subject getting rid of the accuracy requirements which I think is still a massive mission creep as far as this work group is concerned.

KEN RENARD: Okay. All right. Okay. I saw Peter DeVries, I think on this. If you have any comments about that.

PETER DEVRIES: No. I don't have any comments. I don't have any comments. I actually think that's probably how I intended to use placement because mostly I refer to it in network placement and routing so I think we're on the same page there and it was just difference of terms. So, agreed.

KEN RENARD: Okay. And we think we could get rid of the accuracy requirement as well?

PETER DEVRIES: Yeah, I'm fine with it. I think we hashed through that at the last session. I had no objections.

KEN RENARD: Given that RSSAC47 and the requirements and hopefully the underlying ground truth of the entire RSS is that accuracy is extremely important. Hopefully, we can assume that while there are some—so, we're measuring the RSS, we're not measuring anybody else that's pretending to be the RSS. Getting a little bit into rogue talk here but, so, how about this? I could propose some new text for that and we can discuss that via email and then hopefully conclude something at the next core writers group. Okay. Just taking some notes here. Okay. Excellent. So, the next thing to talk about is section three, the measurement section. Again, this is where we are aggregating the measurement requirements from each of these user narratives into something like pseudocode to describe what this tool should do. And Andrew has made some significant contributions to the document here on this. And if you could talk to this section, Andrew.

ANDREW MCCONACHIE: Yeah, sure. Thanks, Ken. So, yeah, as Ken was saying, the idea here with this measurement section is to get down into the nitty gritty of what precisely the tool will be required to do in order to accomplish what's stated above in section two in the narratives. So, this is going to be kind of the highly technical section of the document. And so there's two measurements I've written up here, DNS query latency and also traceroute. We'll start with DNS query latency first. But I guess before that, Fred, it looks like you have a bit of a meta comment here on the section before we get into DNS query latency. Do you want to speak to your comment?

FRED BAKER: Sure. What I'm looking at is that a query latency or a traceroute inherently uses the protocol in question before UDP, whatever and it goes point to point. So, if this is coming from a RIPE Atlas probe, that may or may not be the path that is being used by the user. If we're using UDP to do the traceroute or whatever, that isn't going to tell us much about TCP so it doesn't tell us how to optimize that. Just looking at these two items and wondering if they belonged in a different section. And you'll find as we go through the document, I actually put them in, you know, suggested a different placement for them.

ANDREW MCCONACHIE: Yeah. I saw that you suggested them down into the existing tools.

FRED BAKER: Yeah. Which ping and traceroute are existing tools. That was my logic for putting them there.

ANDREW MCCONACHIE: Yeah, they are existing tools and they can certainly be used by anyone to do a certain amount of debugging on their own. I imagine section three to kind of be talking about how specifically the tool would use something like DNS queries or traceroute or ping to capture specific things. So, like to capture things that are common across all runnings of the tool, if that makes any sense. But I think you have a point in that they should also be talked about here in section four because they are existing tools and people use them all the time.

KEN RENARD: I understand, Fred, what you're getting at. The other existing tools perf_root, the Visualizer and Verfploeter, those seem like higher order tools. I mean, ping and traceroute are hopefully ubiquitous enough that they're really being referred to as building blocks for this tool. And in reality, what we're describing here that this tool would do to measure the local perspective, we're describing it in terms of traceroute, in terms of Dig. There's no requirement that this tool actually use those commands but it probably just makes sense to use those as building blocks.

FRED BAKER: Well, and I would agree that they're their building blocks, they're words we use to describe a type of measurement and type of measurement might be done with the TCP SYN for all we know.

ANDREW MCCONACHIE: So, how about this? I'll take an action item to write up ping and traceroute here under existing tools and I guess you're not talking about Dig so much. Although we could include Dig in existing tools as well, if we're going to include ping and traceroute.

FRED BAKER: Yeah, we could probably add a section on Dig there. Or if you want to resolve my comment there in section three.

ANDREW MCCONACHIE: Let me just take a couple of notes here.

KEN RENARD: Actually, Andrew, the only reference to ping is in section 4.4 header. I didn't find anything else where we're actually using ping.

ANDREW MCCONACHIE: 4.4 header. Wait. You mean the 3.1?

KEN RENARD: I only see two instances of the string “ping” in the document.

ANDREW MCCONACHIE: And where is that instance?

KEN RENARD: They're both on page nine at the bottom in that 4.4 header and then in the footnote for number six on mapping. So, I don't think we actually use ping anywhere in our description of what this tool does.

ANDREW MCCONACHIE: We may say ICMP somewhere like here in Verfploeter. Yeah, I don't know.

KEN RENARD: You know, other than ping is another way of potentially getting round trip times.

ANDREW MCCONACHIE: You're right. We don't say ICMP here in section three. Okay. Well, so afraid if I do that write-up in section four on those three tools, traceroute, ping and Dig, would that resolve your comment here?

FRED BAKER: Yeah, I think so.

ANDREW MCCONACHIE: Okay, cool. Okay, then let's talk about DNS query latency. So, in the writing group, we had a discussion about the kinds of queries that we would want this tool to send. I think we focused on initially—or initially

I had written this up to query random NS records for random TLDs and that was the only query. And then in the core writing group, we came up with the idea of using hostname.bind queries and then also querying the .com name servers and then also querying the DS record for .com and the details of that are really here. And so based on that discussion, I rewrote section 3.1 last week and this is what we're left with. Ken, do you want to talk about your comments? I don't see any other hands raised. If anyone has any comments now. Okay. I see you're un-muting. Go ahead.

KEN RENARD: Okay. thanks. So, just one thing, I agree with the highlighted text here about the end of second paragraph for TCP, the time measure would be the difference in timestamps between the SYN to the FIN. I'm just wondering if we know what does Dig report, hoping to leverage the use of the dig tool, it would be nice if these things synchronized.

RAY BELLIS: That's right here. I don't know, I can try to find out [given] who we work with.

KEN RENARD: Yeah. And the details, if we're talking about the, yeah—I am not sure exactly what we're going for here. The measurement. The important thing is, how long does it take in a sense that we can compare it to others. If it takes an extra couple of milliseconds to shut down the internal TCP session or whatever, it's probably not too much of a big

deal. But the other question I had was about so we're doing 10 queries for each. So, for the use cases we describe here, the question begs, do we want to run this tool hourly? Do we want to run this tool for indefinitely like an Atlas probe? Do we want to run this tool just for a short period, maybe a week or a couple of days so that I can find out what my local perspective is, find out if this is underserved, answer a question and then stop? So, all those ideas are out to the group for comment but my main question here was, are we going to run these 10 queries back to back for each root server identifier right back to back? Or do we want to wait a few minutes? Do we want to do it over an hour or a day period for these 10 or? Any thoughts on that? And I see Brad's hand up.

BRAD VERD: Yeah. Again, I go back to the intent when we started this work. The intent here, again, was to try to help answer the person who—one of two things. One who said, "I'm in an underserved area" or two said, "I want to host an instance of a root server." So, then you could provide them tools to provide that data to the people who would be providing that, one, answer the first question, if they are underserved and provide that data to the RSOs and two, if an RSO was interested in putting an instance there, they would need a bunch of data. Anyways, none of this was meant to be run in at least the intent when this started, it was not meant to be a new monitoring system that would happen in perpetuity. It certainly could be or maybe if somebody is looking for an instance, they run this whatever this set of tools, not just once but for a period of time and provide that data. But that was the intent when we started

this. And again, if it's changed that's fine, I'm just going to remind people.

KEN RENARD: I think that's good and I think that really should be said somewhere in the document, probably up near the top or possibly in a section. We often compare this to Atlas probes which should be run for a long time. So, this tool is not meant to do that necessarily. I'll take that as an action item to describe the time period of how this tool would be used.

FRED BAKER: Well, yeah. At least as I understand it, the tool would be used when somebody has a question and they would use it in a way that answers their question hopefully.

KEN RENARD: So, then narrowing in, I definitely see the value in getting 10 measurements because something can be dropped. 10 successive measurements here. As the pseudocode's written, it's basically almost one instance in time. Do we want to say anything about the pseudocode level that this should be run periodically, hourly, daily, whatever, or just leave that specifically up to each user narrative? Andrew?

FRED BAKER: Personally, I would leave it to the user narrative.

ANDREW MCCONACHIE: Let me make a suggestion. We can add another section here under section three called timestamp and we just treat it like another measurement and we just say, "When the tool begins, it takes a timestamp and we can describe UTC or whatever, how we want to record that timestamp and then when the tool finishes, we also record another timestamp." I don't know if we need to say more than that, about how, you know, what kind of regularity with which this tool needs to run. Maybe that just gets left up to the user but I do think we should describe how we want time to be recorded so that if there's multiple instances of this tool, multiple implementations of this tool, they're all recording time in the same manner.

KEN RENARD: Okay. And I'm falling back in my mind to the Dig output which I'm pretty sure gives a timestamp with a time zone in it.

ANDREW MCCONACHIE: Are people okay with me writing another section called timestamp, which is really short, just describing how we record time?

KEN RENARD: That sounds good to me, Andrew. I'll see what I come up with about clarifying what Brad was saying that this is the time period to run this tool is to answer a question and see if those actually fit together.

WES HARDAKER: It sounds like you need the output formats for each of the data types that we expect output rate. So, for timestamps which is different than the latency measurements, timestamps, you might want to reuse the ISO format, for example. Latency measurements, maybe you want to standardize on milliseconds versus seconds, that type of thing but [inaudible]. Every sort of output likely needs a different described standardized format.

ANDREW MCCONACHIE: Correct. And I did that. So, for DNS query latency, I think I said milliseconds. Well, I should say that—last version I said milliseconds.

FRED BAKER: Well, I just did a Dig to an ISC system. I don't see a timestamp in the dig out, but—

RAY BELLIS: Fred, it should be the very last line where it says when but it's just in—it looks something like an [inaudible] format.

FRED BAKER: Yeah, you're right.

KEN RENARD: Now, question for the group. Are we describing what information needs to be collected here and leaving it up to an implementer to come together with specific formats if it's JSON, YAML, whatever? I think

we've somewhat punted on the idea of writing the tool and just more describing what it should collect.

ANDREW MCCONACHIE: Well, we previously decided we weren't going to go into how this should look in JSON or how this should look in YAML. I think that was on a previous meeting. But I do think we need to say this value should be in seconds, this value should be in milliseconds or something like that.

KEN RENARD: Agreed.

ANDREW MCCONACHIE: Okay. Well, conscious of the time, given the fact that we're talking about time, I'm going to continue. I'm going to go back down to DNS query latency. I'm surprised I'm not hearing people telling me that these are the wrong kind of queries to be sending. I expected lots of opinions about the types of data we'll be sending and getting back. Duane, please go ahead.

DUANE WESSELS: I was debating whether to say anything. I was just going to say that that might still be coming from me. I'm currently going through the document that I'm up higher making some comments but I haven't gotten to this point yet.

ANDREW MCCONACHIE: Okay. Well, thanks. Yeah, we definitely appreciate that.

KEN RENARD: The one thing that, Andrew, that I don't see that got carried over from our discussion in the core writing group was the timeout values. To the larger group, what we had wanted to do specifically for the use case of identifying or informing the underserved decision is really increase the timeout value so that we could really get an idea of how long is that round trip time? If we increase the timeout value at say 10 seconds versus a typical like four, that would allow us to really see how underserved the area is. Give us a wider set of timeout values that could also lead us to given that timeout value, a large timeout value, it could then be interpreted if the timeout value was X, you can now interpret your availability because if the response came back in six seconds, when your timeout's four, that's essentially unavailable. Ray?

ANDREW MCCONACHIE: Ray, I think you're muted.

RAY BELLIS: Sorry. Yes, I was. Sorry. For the Dig commands, we do have options that we can actually disable retries, for example. And I would suggest that if they got the situation where a client is having to retry to get the root system, that's actually a bad thing. So, I would say [inaudible] tries, set the timeout to one second. Because anything worse than that as far as I'm concerned is not functional. Yeah, that's not a suitable quality

service and certainly with the Atlas probes, I never see anything exceeding one second.

ANDREW MCCONACHIE: What if all queries fail at one second but they might've passed at say five seconds, is that interesting to know?

RAY BELLIS: To me, I'd still say that was a failed query but, yeah, it depends on whether we're trying to maybe follow up with a ping or a trace. I mean, if you're not getting back an answer within a second, then the system is really not working.

KEN RENARD: Well. I think one of our use cases, especially for underserved was an area that's probably very remote and does not have good service. So, if we had, let's say it's a seven-second round trip time to the best root server identifier and had a 7.2-second round trip time to. com or one of the open resolvers, that's about as good as you're going to get. So, that idea of that you can compare it. If you know that the timeout is five seconds or seven seconds, you can then make your determination of my tolerance for time-out is one, my tolerance for timeout is 0.9. You can infer that but essentially losing any round trip time over N, losing that information, might be less informative.

ANDREW MCCONACHIE: One thing we could consider doing is just start with a high value and then decrement it down to a low value one, right? Like, so you would start with 10 and then halve it, and then go to five and then go to 2.5 and then go to one or something like that.

RAY BELLIS: Yeah. Well, I think I should do fractional time that's as far as I recall anyway so we might try to just one, two, four eight, if you're going to go that route. My own experience is I've never seen a working root server query from an Atlas probe that's anything beyond one second. I have seen as high as 700 milliseconds but no worse. Anybody who's getting worse than that must be [inaudible] something like a satellite link because terrestrial [inaudible] just simply not that high in normal usage, unless there's something severely wrong with local Wi-Fi congestion or something like that.

ANDREW MCCONACHIE: Peter, go ahead.

PETER DEVRIES: Yeah. I will say we do the same thing. We limit it to one second because we found anything else is it's broken at that point. And so, I would also say that I think we're running 130 queries here so we kind of have to be aware of how long it takes. If you set your time out for at five, you could be running for minutes. So just want to take that into consideration to what we set the time out for.

ANDREW MCCONACHIE: Duane?

DUANE WESSELS: I just wanted to note it in case anyone remembers RSSAC047, we set the timeout value to four seconds there which is pretty high. I think one is sufficient but just as a data point, that other work used four seconds.

ANDREW MCCONACHIE: So, what I'm hearing is we should maybe just set this to one. Is that okay? Is everyone okay with that change?

RAY BELLIS: It's fine with me certainly. The other thing I'll point out is that from the recursive approach point of view anyway, any full featured recursive resolver is generally going to use the fastest root servers they can see. So, yeah, if you've got one that's really, really slow or even several that's slow but you've got really fast access to the others, those others won't get used.

ANDREW MCCONACHIE: Is that an old hand, Peter?

PETER DEVRIES: It is. I apologize.

ANDREW MCCONACHIE: No worries. Okay. So, we'll set that to one then. Any other comments on these queries or the pseudocode? I don't know if there's a standard for pseudocode or whatever. I just made up some pseudocode, hopefully people can understand what that means. But any more comments on 3.1 about anything about 3.1? Okay. Not seeing any, we have about five minutes left. Maybe we can talk a little bit about 3.2. Ray, do you want to talk about your comment here, the UDP traceroutes? Reading this, it looks like you're just saying we should also be doing TCP traceroutes which seems reasonable. We could do both. Is that an idea? Should we be doing both UDP and TCP traceroutes? Ken?

KEN RENARD: I thought the purpose here was to get the path to try and inform a decision on routing optimizations. Does TCP traceroute add anything?

WES HARDAKER: It might if there's load balancing involved.

KEN RENARD: What was that, Wes?

ANDREW MCCONACHIE: Say that again, Wes.

WES HARDAKER: I said it might if there's like some sort of load balancing involved where things are sent one way versus the other. But that would typically be at

the end system, not at the middle of the path, or at least I hope that's not happening in too many places. I don't know. Maybe SDR is changing all that and it's going to happen in a lot of places.

ANDREW MCCONACHIE: So, there might be some value in having TCP traceroutes if there's a chance that UDP and TCP are taking different routes?

WES HARDAKER: I mean it would be cheap to execute and good to report, and so why not do it? It's not like it's going to add—that's not a huge overhead, right? To do both?

ANDREW MCCONACHIE: No. And it can be done at the same time. So, theoretically it shouldn't take the script or whatever any longer to execute. So, let me add—I'll just respond to Fred's comments. Say add [traceroutes]. Okay. There was some discussion in the core writing group about setting a maximum TTL of 32 hops because there's so many intermediary gateways on the Internet that don't respond with that ICMP packet that's the drop notification of the probe. Another option is something I do in perf_root which is, I set a maximum TTL but if I get five non-responses, I just halt sending probes. So, we could specify doing both because having a maximum TTL of 32 means that if you encounter a bunch of gateways that don't respond to you, then you're still going to go all the way up to 32 even if like the last 20 of them don't respond to you, which can take a very long time. Ken?

KEN RENARD: Yeah. So, is there any value of waiting that long extra time for a response that's probably not going to come back anyway?

ANDREW MCCONACHIE: I don't think so. If anything, it's just going to annoy the people running the script. I will incorporate this idea and then people can comment on it in a later version. And I think we should probably end here because my clock is showing about 30 seconds left. So, thanks everyone. Ken, I'm going to turn it back over to you.

KEN RENARD: Thanks again. I encourage everyone to contribute to the document. You can participate over RSSAC caucus mail list or directly in this document. And I don't know, Ozan, if we have another meeting scheduled but it should be near the middle to end of April. And with that, I also wanted to turn it back to Abdulkarim to see if he has anything final to say. And I want to thank everyone.

ABULKARIM OLOYEDE: Thank you very much, Ken. I don't think I have anything else to add. Yes, I think Ozan [inaudible] and thanks [inaudible].

OZAN SAHIN: Thank you, Abdulkarim and Ken, and thanks everyone for participating today. For the participants that are non-RSSAC caucus members, I just

pasted the links to the RSSAC webpage, as well as the RSSAC caucus webpage if you'd like to learn more about these groups. So, thanks again and have a great rest of your days. Could we please stop the recording?

[END OF TRANSCRIPTION]