

---

OZAN SAHIN: Good morning, good afternoon, good evening. This is the RSSAC caucus RSS metrics work party call held on the 12th of June 2019 at 15:00 UTC.

On the call today, we have Duane Wessels, Abdulkarim Oloyede, [inaudible], Jaap Akkerhuis, Jeff Osborn, Kazunori Fujiwara, Ken Renard, Matt Larson, Ryan Stephenson, and I see a nickname, [inaudible], so if you could announce your full name, we'd appreciate that, of the roll call.

KEVIN JONES: That's for Kevin Jones.

OZAN SAHIN: Thank you, Kevin. Sorry for missing that. And from support staff, we have Andrew McConachie, Steve Sheng, and myself, Ozan Sahin. I'd like to remind you all to please state your names before speaking for transcription purposes, and mute your microphones when not speaking. Thank you, and over to you, Duane.

DUANE WESSELS: Alright. Thanks very much, Ozan. So Russ Mundy, the work party co-chair is not on the call yet, but I think he'll be joining us later. I guess my idea for today was to sort of go through some of the topics that I had sent in the e-mail from yesterday. Let me bring that up real quick.

Yeah, so I guess first of all, if you're on the call here and haven't been watching the document very closely, please take a look at the Google

---

*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.*

document and read through it. There have been some changes since the last time. Not all the changes probably are visible as changes, because I think Steve has been going through and accepting some of the edits that are maybe appropriate edits to accept. So you may want to sort of take a read through it and see for yourself what has changed.

Sent an e-mail out yesterday highlighting a few of the things that I've been working on, and I think we'll just sort of go through those one by one, and hopefully have a discussion around those.

One of the first topics is – and there are still some comments in the document about these definitions. We have some definitions, some terminology which is copied from another RSSAC document, and those are not really in contention at this point, but there are some comments about our use of the words “measurement” versus “metric.”

For the most part, we're sort of following a suggestion, I guess from Paul Hoffman originally, and hopefully I didn't misrepresent Paul's thinking too much, but the idea here is that our use of the word “measurement” kind of refers to an individual measurement or individual measurements and “metric” refers to taking a collection of measurements and then sort of applying some math to that and then coming up with sort of an aggregated metric over some period of time.

So there's some different opinions on whether or not these terms are appropriate. I feel like at this point, we don't really have better definitions, and so those definitions sort of stand for now, and I think they're used consistently in the document. But I'm certainly open to

---

other suggestions for how to improve this. Would anyone like to discuss this point on the call at this time?

Not seeing any hands, there was a response from Kazunori about maybe let's not redefine terms that are already well defined in English, so we should refer to these as RSS and RSS measurement rather than standalone metric and measurement.

So we can take that as good advice and try to apply it throughout the document. I'm a little bit concerned it may make the document a little bit awkward to read, but I'm willing to give it a try.

I see there's some comments in the chat about this, which is fine. But if no one wants to speak to this topic, then we can move on to the next one. Okay.

Kind of in a similar vein, we had some comments in the document about the use of the initialism RSO in the first half of these metrics, it previously talked a lot about measurements against a single RSO. My concern was that we were confusing the terminology for a root server operator with the service that was provided by that operator, and particularly in the case of Verisign who operates two root servers, two root letters, if you will, that was a little bit weird.

So at this point, most of the occurrences of RSO have just been replaced with root server, and I guess in section four, you'll see that it mostly says things like root server latency or root server correctness rather than RSO latency or RSO correctness. Any comments about this particular change or topic? Alright, doesn't sound like it.

---

Let's move on then. So the next thing in my e-mail was about measurement intervals, and I think from the time of our last call, some of these metrics had measurement intervals that were defined at, say, one minute – the idea was you take a measurement every minute, in some cases for other measurements, it goes five minutes, and then I know for at least one of the RSS metrics, every hour was the measurement interval.

And I went back and looked at the meeting transcript from the April meeting where we presented this work to RSSAC, and specifically at what the discussion around the intervals was there. I think that for the most part, people felt that one minute was maybe on the high side of too frequent, and the five minutes would be appropriate.

I also feel that there's some utility and value in having all of the measurements done at the same interval that makes it easier for us to think about them and maybe even to implement them.

So at this point, the document has been modified a little bit so that – I think almost all of the metrics say that they should be taken at a five-minute interval, which is what I would propose going forward. Even if there are cases where that's too much data, we could later change it so that – even though you maybe make a measurement at a frequent five-minute interval, you can ignore some of the data that you collect.

Hi, Russ. I see that you've joined. I'm just in the middle of going through some of the topics in this e-mail that I sent out yesterday.

---

RUSS MUNDY: Yes. excellent. Thank you.

DUANE WESSELS: We're just in the middle of talking about measurement interval, and I think I've said all that I want to say, and I would like to open it up for input from others if they have opinions on this. Clearly, we're missing our good friend Paul who often likes to talk to me about these things. But hopefully Paul's reading his e-mail and will respond when he has time.

RUSS MUNDY: I'm not sure if we're trying to do hands today or not, but let me just pipe in. I think the five-minute interval is an excellent point for us to land on at this time, and if we do need to make variations on that, we can do it later. But five minutes, I think, is a good point for us to use here.

DUANE WESSELS: Okay. Thanks, Russ. Alright. Yeah, I'm not seeing hands, and please, if anyone wants to speak, just go ahead and speak up. And if we get to the point where we have too many people that try to talk, then I guess we'll use the hand system.

Okay, so let's move on to the next topic, which was PTI naming performance reports. So as I said in the e-mail, during the April meeting with RSSAC, these reports were sort of held up as an example of something similar to what we're trying to do here, that is done very nicely, and I'm inclined to agree. These reports are very well done and pretty clear to read, I think.

---

The one thing that's different between the way that the PTI IANA performance reports work and the way that the RSSAC metrics might work is for the most part, the PTI metrics are sort of – I would say they're self-reported. The systems that PTI uses have been modified to record and report the times that certain operations take, and almost all of the things that they are measured against are – they're driven by activities where TLD managers submit changes for example, or they're driven by work coming to them in the form of TLD changes or investigations and things.

So that's different than what we're considering for the root server system where for the most part, we've been talking about doing what I would call active measurements, maybe by a third party, to measure system performance of servers and things like that. So that's a little bit different, but I do like the way that the IANA reports look.

So for some of the metrics in the document, like the one that you can see on the screen here, right at the bottom, there is a couple of tables. The first table says example results, and this is sort of the older format that I had in mind, and you've probably seen this before. The one below it says example based on PTI CSC report, and you can see it's a little bit different. Ozan, maybe you can scroll down. I think maybe the table is broken.

Yeah, so there's the rest of it. It's just presented a little bit differently, and one thing that I liked in the PTI report is that in the cases where the metric or the threshold level was met or satisfied, it's highlighted in green, so that's shown here as well. In cases where it's not, then it

---

would be highlighted in a different color, I think yellow or something like that.

So there's this example, and then I think maybe is there another example down at the next one too, Ozan? Can you scroll down to the latency one? So this one is maybe even a little bit more interesting, because kind of what's being proposed here is a way to show the range and accuracy of the measurements, is to include all these percentiles. So for example, this one table is just about UDP, IPv4 latency on a specific date. It shows the latency of 50 milliseconds, and shows a hypothetical threshold value of 100 milliseconds at the 95th percentile, which is what that threshold percentile shows.

And then it has some more information about the distribution of measurements, so it shows that there were 280 measurements included in this calculation, and there are various percentiles from 5th percentile, 25th, 50th, 75th, and 95th.

So you'll see that the 95th percentile latency of 50 milliseconds matches the one at the top, the green line. So that's sort of how the PTI reports looked, and I felt that was pretty good. Any comments or feedback on adopting this approach for the RSSAC metrics?

Well, not hearing any objections, I guess at least one thing that we can do then is continue to put these style of examples in some of the remaining metrics, because that hasn't been done yet. So [inaudible] that for future work.

Okay, moving on then, I'll see if I can find something really controversial to talk about. So far in the RSSAC metrics work, we haven't spent a lot of

---

---

time talking about the location of the measurement probes, although I would say there is this assumption that the measurement probes would be widely distributed, and I would say sort of far away or distant from the root servers that they are intended to measure.

So in-between the probes and the servers, you might have any number of third-party networks involved, and that makes certain things a little bit complicated. One complication arises from when we talk about doing correctness measurements and we have a very high bar for correctness, for the most part I would say we're targeting 100% correctness, but if we have these third-party networks in there where nefarious things can happen, where packets can be spoofed or fiddled with or dropped, then it gets sort of tricky, because we are essentially holding the root server operators responsible for things that are out of their control. And I think there's a similar issue when we're talking about latency.

Today, we've been, again, suggesting that latency might be measured from these distant probes and going over networks whose utilization and congestion are out of the operator's control, and is it fair to attribute latency in those third-party networks to the operators?

So one way to address these complications, I guess, is to use lots and lots of measurements. The more measurements that you make, then you can sort of eliminate some of these outliers. If you have concerns about things like spoofing and interference, then you might have to do things like look at a whole series of measurements or perform measurements to other targets and see how they compare. It would

---

maybe eliminate measurements from probes that you suspect are not reliable.

So that's one approach. Another idea that I've had and had some discuss with others is maybe for at least some of the measurements, we should consider having probes that are closer to the servers themselves to eliminate some of those unknowns and eliminate some potential problems from third-party networks.

I know in the case of Verisign actually, some of the measurements that we do, the probes are located actually quite close to the servers, and I think it might be the case for some of the new gTLD SLA measurements as well, although I'm not quite as sure about that. Does anyone else on the call know more about how those measurements actually work for the gTLDs?

When I was looking through some of the – I guess it was the applicant guidebook, or one of those supporting documents, I remember reading something that some of the probes are located close to the servers.

RUSS MUNDY:

Duane, I don't know from a factual basis, but from the description that was provided in the set of specs prior to them actually doing that, the requirement was for the testing points, the probe locations to be available on, I think it was five of the continents that were identified, or all the continents except Antarctica, but the one thing that was not clear to me from when they did that, if they were intentionally planning on probing far away or if they were intentionally planning to probe close by.

---

So I know there was a dispersal requirement. I don't know how they intended to use it though.

DUANE WESSELS: Okay.

JAAP AKKERHUIS: I have some notes about technical requirements for the new TLDs, which are more specific than what's in the guidebook. But I guess the best one to ask are the Swedish people, because they actually did tests for the incoming gTLD, so they might know best what they did.

DUANE WESSELS: Okay. Thanks, Jaap. Maybe you can put us in touch with individuals from the Swedish continent.

JAAP AKKERHUIS: I think it's just [inaudible] but I'm not sure. Remind me if I forget. I'll try.

DUANE WESSELS: I guess what I'm sort of proposing or maybe leaning to is that maybe for some of the measurements we want, we'll have a set of probes that are close, and for other measurements, we want to have a set of probes that are not close.

I don't like that, because I think it's unfortunate that we would have to have two sets of probes with different purposes, but from the point of

---

view that if we're able to design something from scratch and design it the way that we think it really should work, then that sort of feels like the right thing to do to me.

RUSS MUNDY:

It seems like that is an answer, is two separate sets of probes, but depending on the construction of the probes and how they're managed, a given probe could be testing things nearby as well as doing other tests for entities further away from a network topology perspective. So we don't, I don't think, have to necessarily have a duplicate set of probes, just structure the test in such a way that you can do both types of tests from a probe.

DUANE WESSELS:

Yeah, I think that's true. and I think some of this, we really probably need some good input from the operators themselves, the extent to which they would – how close could the probes actually be in the same rack for example or something like that? And would such probes in such racks be able to reach out to the Internet as a whole and do measurements? Without that level of cooperation, then we may be more limited in what we can do.

I see in the chat, Ken, you had an idea about – is it to propose using sort of passive measurements to detect spoofing?

RYAN STEPHENSON:

Hey Duane, sorry to interrupt. I just wanted to clarify if what I heard was correct, meaning that the RSOs would have probes that are near to –

---

maybe within the same rack or something, or within a nearby facility to the RSOs' locations? I apologize.

DUANE WESSELS: Well, that is sort of what I'm considering, yeah.

RYAN STEPHENSON: Okay. I just wanted to make sure I understood that. Thank you.

DUANE WESSELS: So again, if we're concerned about networks beyond our control interfering with our measurements, then the closer that we can get to a root server for those measurements, then the more confident we can be in the results.

And certainly, if we can get within the same rack, then that's very good. If it's not within the same rack, if it's within the same data center or the same whatever, then that's still good, but maybe there's still some uncertainty there. I think that's what we need to talk more about.

RUSS MUNDY: Yeah. I fully agree. We need to have more discussions on this, because the concern that is in my head about having them having the testing done physically very near the root server machinery itself, it does not, I don't think, get to the question that a lot of people, when they ask about the root server system as a whole are worried about that does include these other problem spaces, or challenges if you will, and it puts

---

us in kind of a quandary of a position that if we flavor it or if we have a scheme that only looks when you're real close to the machine, versus from some spots out around the network, I think that would not produce a result that would be well-received in the wide community. Let me just put it that way.

I'd love to hear other people's thoughts on this too.

RYAN STEPHENSON:

Sorry to interrupt, Russ, and again, Duane, sorry for going off topic here. So to just kind of reiterate, some discussion between having the probes as close to the root servers as we can get, but also because of the fact it can look kind of that the data is necessarily kind of stacked to be favorable to root servers because it is close enough that the community may say, "But it's not a true degree of measurement maybe from a location that the root server is not nearby," in very layman's terms. Is that kind of what I'm gathering?

RUSS MUNDY:

Yeah, that's really kind of the concern I was trying to express, Ryan.

RYAN STEPHENSON:

Understand. Okay, thank you. Sorry, Duane, for interrupting. Go ahead.

DUANE WESSELS:

No, you're not interrupting. This is discussion we want to have. So this is all good. There's a discussion going on in the chat as well, and I'm trying

---

to sort of keep up with it and follow it. Ken, are you able to speak on the call? Or are you only on the chat?

KEN RENARD:

Okay. Yeah, my thought was that putting the measurements close to the RS, root server, is really just for validation. If a packet or a response is rewritten in root to the probe, we would at least be able to verify. I definitely agree that putting true measurements or a probe close to the root servers does skew it in our favor, so maybe not take the latency or availability metrics from that, only use that as a way to debug whether metrics from far off lands are being intercepted, rewritten. Just a validation of the metrics.

DUANE WESSELS:

I see. That's interesting. We should probably – I'm sensing myself getting into the implementation details, which is probably not the right thing to do at this point, but I think that's an interesting idea. So we can talk more about that. But I see what you're saying. Okay.

RUSS MUNDY:

Ken, if I could, are you then suggesting that it might be appropriate to have different placement for probes with respect to a particular measurement?

KEN RENARD:

Not necessarily. I'm thinking more when we talk about correctness, we expect 100%. Sometimes, responses do get rewritten, going back to,

---

say, Atlas probes, we could look at – when we don't see 100% correctness, we could at least investigate what's going on. if it's the probe that received the incorrect response, and I see that [packet capture] or something at my root server that said, "Hey, I saw that request, I saw the response, it went out with this value, but it was received as a different value." That's just a way to validate the response or validate the measurement. Is a probe measuring what the root server responded with, or what it received from somewhere?

DUANE WESSELS:

So Ken, my initial thought was that by putting some probes, or a subset of probes close to root servers, you can have more confidence in the correctness measurements.

KEN RENARD:

Agreed.

DUANE WESSELS:

What I sort of hear you suggesting is – another thing you can do if you have these probes or some kind of devices close to them is that you can compare what it looks like close and what it looks like far away, and this allows you to detect cases where things are being spoofed.

KEN RENARD:

Exactly.

---

DUANE WESSELS: Which I think is very interesting. I'm a little less sure the extent to which detecting interference is part of root server system metrics. Maybe it is, maybe that's something that's very useful for people to know. They want to know the extent to which spoofing is happening from root servers, and so that could be done that way.

KEN RENARD: It's kind of an attempt to say, "These are things that are beyond the root server operator's control," the network path or the interference or spoofing, just be able to identify that and be able to validate those responses, validate the metrics.

DUANE WESSELS: Right.

WARREN KUMARI: It seems that having probes close to the server and far from the server, both give useful results, but they give very different sets of results. A probe close to the server gives an understanding of the performance or latency of that particular server, and then further away, gives you an understanding of what the user would be seeing.

And I personally think it would be worth having both of those. You could even – if you wanted to get fancy, subtract the average latency or performance of the probe near the server from the network so that you could get a better understanding of how an increase in peering or an increase in locations would improve, or not, the user experience. Was that in any way coherent, or did I just go on a complete ranty tangent?

DUANE WESSELS: No, that's it, that's what we're getting at. So I appreciate the feedback. It matches my understanding as well. I think this has been a good discussion. I think that going forward, we'll give this a strong consideration in the document to suggest that maybe there's – if not two sets of probes, at least probes at different locations for these different purposes, and as Ken and Warren have pointed out, it sort of allows us to do even more things than maybe we originally had in mind.

RUSS MUNDY: One thought that crossed my mind as you were just describing that, Duane, is that perhaps instead of describing the devices that are doing the testing as probes, which kind of infers that they're probably a single thing, would it make sense perhaps to talk about the test points for each of the respective metrics? Because they may need to be different, and then let the actual implementation later on decide if it's one physical probe doing multiple tests or not. Does that make any sense at all?

DUANE WESSELS: I think that's a fine approach, and I think we can certainly work with that. One of the things that I feel like we struggle a little bit in this work party is exactly how descriptive or prescriptive we want to be with the probes. Do we want to leave it sort of very undefined [and it's sort of an] implementation detail, or do we want to be more specific?

---

One of my goals throughout a lot of this has been to define measurements in such a way that they can be done from an existing platform like RIPE Atlas. What we're talking about now with close probes and far probes, that's not possible, strictly from RIPE Atlas only. But I think that's okay.

WARREN KUMARI: An option might be – and this is probably a bad idea – to have individual letters, or what we're referring to this week as, do tests to other ones. You could ask potentially root server operators to do tests to other letters and report those statistics as well.

DUANE WESSELS: Yeah. I think that's a reasonable suggestion. There are some details to be worked out, but I think that's a reasonable thing to consider.

WARREN KUMARI: It also depends on if you're asking people to do a test to each one of everybody's locations, or just to the Anycast address. Those are two very different sets of requirements on people and two very different sets of metrics, obviously.

DUANE WESSELS: Yeah. And to be clear, Warren, so far in this work, nothing that we're suggesting would try to do measurements against, say, all of the instances of an operator. It's just you send to the Anycast service

---

address, and whichever one you hit, you hit. That's what we've done so far.

WARREN KUMARI: Thank you, and apologies for not knowing that.

DUANE WESSELS: That's okay.

RUSS MUNDY: The other thing that I think the work party needs to decide, particularly with respect to correctness, is what is the right set of things for the metrics party to define as part of metrics, and any action associated with that, and what would be – I'll describe as a trouble report or troubleshooting, as Ken was talking about earlier when a response left the root server and it was correct but it got to a location that was doing testing and it was incorrect. Is that a network troubleshooting thing that's not really a metric, or is that something that is considered part of the metrics?

And I don't have a strong feeling one way or the other, but I think it's something that occurs, and the work party needs to figure out what should be said about it.

DUANE WESSELS: Yeah. Russ, for me in that case, I think we could develop a metric around that, but in my mind, it wouldn't be a metric that is sort of tied

---

to or at least could be used to – I wouldn't blame an RSO for something that happened in that way. That's something out of their control.

RUSS MUNDY:

Right, but unless we have some mechanism in place that we describe in the metric, along the line of what Ken was saying, to do some kind of comparison, it wouldn't necessarily be obvious why the location that was running the test case got an incorrect answer.

It could have been from the RSO itself, or it could have been somewhere along the path. And I don't know how much we as a metric work party ought to try to identify how one would go about figuring that problem out.

DUANE WESSELS:

Yeah, that needs more work and discussion, I think. This discussion has been good. I think maybe in the interest of time, let's move on down to another – I guess there's only one more topic. So, alright, have we exhausted the discussion about probe locations, or would anyone like to say more?

WARREN KUMARI:

I apologize again for joining late, but what exactly is this metric supposed to show? Is it supposed to show the experience of users in general, or experience of users of the root server letter as measured from somewhere? Because these are very different things.

The experience of the users in general requires, I think, measuring from a large number of locations, using a distributed measurement platform like Atlas or something similar. Otherwise, we're going to miss people in Africa or somewhere where we don't have a lot of measurement locations.

If it's to show that individual letters are performing as well as they can, as measured from somewhere close to them, that's a very different set of metrics and numbers and means a very different thing.

DUANE WESSELS:

Yeah, Warren, I wouldn't say that it's perfectly well defined, but I think for the most part, we're talking about the latter, we're talking about metrics against an RSO. How does the root server behave or perform, and less so about the experience of the user.

Now, if we can imagine that we have these probes in different locations, some close and some far, then we can do some of that latter stuff. We can also report some measurements and metrics on user experiences. But I would say the extent to which part of this work is being driven by RSSAC 037 and probably some future service-level whatevers, those would really be tied to the metrics against the root servers themselves and not the user experience.

WARREN KUMARI:

So if I put three really fast servers right next to where the probes are, I will look really good. The fact that that doesn't actually help the user seems to be missed in that.

---

DUANE WESSELS: Well, that's why I think that we would consider both sets of probes.

WARREN KUMARI: Okay.

DUANE WESSELS: I think that if you have a probe close to a server, you can measure that server's response time, we're talking milliseconds, and then if you have probes that are farther away – some definition of farther – then you can do measurements that sort of show the extent to which the root server is distributed and serving different locations, and things like that.

WARREN KUMARI: I guess what I'm saying is – and I guess I worded this poorly – the far away probes, or the total set of probes I think needs to be large in order to provide any sort of representative view of how this is actually working.

DUANE WESSELS: I would agree that that would be great, but that becomes a money problem at some point, right? How much do you want to spend on –

WARREN KUMARI: or use an existing distributed system like Atlas [inaudible] or ThousandEyes.

DUANE WESSELS: Or use an existing system.

WARREN KUMARI: Yeah. Okay.

DUANE WESSELS: Yeah.

RUSS MUNDY: And also, Warren, it might be worthwhile to take a look at the charter for the work party. I can put that in the chat room here. But in fact, the charter says we're supposed to be looking at metrics and defining metrics for both the overall system, the RSS, and metrics for RSOs. There is definitely a stated relationship between the RSSAC 37 and what we're supposed to be doing.

We sort of got shot with a firehose with this thing. I see Ryan already put it there. Good. Thanks, Ryan.

DUANE WESSELS: Alright. I will move on to the last topic from this e-mail, which maybe is pretty noncontentious. In some of the measurements, I'd say these are probably all in the RSS section. We wanted to have some measurements made via recursive nameservers, and wanted to sort of do our best to ensure that we're doing queries for names that are not in the cache.

There were two different purposes for this. That was one purpose, was to do queries, which would be cache misses, and then in the correctness measurements, there was a desire to do queries, and do DNSSEC validation of the responses. And in order to get wide coverage of all the different types of responses that you could get, some of those queries should end up being for names that don't exist, so you get NXDOMAIN responses and NSEC records and things like that.

So that is all to say that currently in the document, there are two types of those queries that are NXDOMAINs or cache misses. In the first one, it is written that you would use names that end with dot-localhost or dot-local or dot-invalid, because those are reserved top-level domains that would never be delegated in the zone.

And then later on, I had this idea that you could just generate algorithmically random domain names that almost certainly would never be real TLDs, like a single letter followed by ten random digits. Essentially, I'm just asking for feedback on these algorithms. Probably, I would propose just choosing one over the other, probably the randomly generated name rather than [inaudible] use the reserved TLDs. I think that would be simpler.

So, any thoughts on that? I would welcome those. Otherwise, absent any feedback, we'll probably just use the random NXD one.

WARREN KUMARI:

I guess i should mention that that does somewhat change what set of probes we should be or infrastructure we could be using. It requires things that definitely do their own DNS lookup and never pass through a

---

---

recursive, which is almost definitely what we want anyway, but that fully puts us into that.

DUANE WESSELS:

Actually, Ozan, can you scroll down to section five, one of those in section five? Yeah, this one. So we're in the section called root server system availability metric, and there are two methods proposed for getting this metric. One of them is that you just use the measurements that were done when you measured the individual root servers, and you just accumulate all those measurements and then do math on them. That's sort of off the screen.

This second one, method number two, says that you can have a probe that has its own recursive nameserver and you could send that recursive nameserver queries. And if the goal is to measure availability, then you probably want to make sure that the recursive nameserver is actually going out and hitting one of the root servers. So you send it a query for something that's not in the cache.

WARREN KUMARI:

Sorry, that was where I was getting. How do we do not in the cache? Keeping in mind NX aggressive NSEC.

DUANE WESSELS:

With aggressive NSEC, you really can't.

---

WARREN KUMARI: Okay. Cool. I think that's where I was going.

DUANE WESSELS: You can do the best you can, I guess, ignoring aggressive NSEC for the moment. If aggressive NSEC becomes more and more deployed, which it probably will, then this method becomes less reliable, I would say. And I don't know any way around that. I don't think there's any sets of queries that we could generate that would always be cache misses in the face of aggressive caching.

So that's something that we should consider, I guess, and maybe we'll decide that this method number two of using recursive nameservers is just not a good method.

WARREN KUMARI: Okay.

DANIEL MIGAULT: One clarification. It's not clear to me, if we are not caching those, how would aggressive NSEC work?

DUANE WESSELS: I had a little bit of a hard time hearing. I heard the first part about –

DANIEL MIGAULT: Is that better?

---

DUANE WESSELS: Yeah, a little better.

DANIEL MIGAULT: So, if we're not doing the cache, I'm confused, because if we're not caching the response, we will not have this [inaudible] cache anywhere.

DUANE WESSELS: The idea, Daniel, is that there is a recursive nameserver that is doing caching, and so the reason that we're considering this as a measurement method is because recursive nameservers, they're designed to choose an appropriate server to talk to, the best one or the fastest one, or whatever, and so that more mimics the real world cases where traffic is flowing through recursive nameservers. But [inaudible], there are tradeoffs here, so I'm not sure it's an ideal approach.

DANIEL MIGAULT: Oh, yeah. I see. Okay, so yeah, it's not just a resolution – it's not [anti cache] [inaudible]. The resolver is being set up [inaudible].

DUANE WESSELS: Yeah.

DANIEL MIGAULT: Okay.

---

DUANE WESSELS: So we've got just a few minutes left, and I wanted to, I guess, remind everyone that I think the next time that we're talking about this work will be with the entire RSSAC at the ICANN meeting in Marrakech, which is in about two weeks. So as before, it is my understanding that those sessions will be open to caucus members, so everyone is welcome to join into that.

I think given that the ICANN meeting is in Marrakech, the time zone may be a little bit difficult for some people, but please look for the announcements of those meetings and please join us if you can. That's correct, Steve, yeah? Anything to add about the Marrakech meeting?

STEVE SHENG: No, Duane. Ozan already put in the chat the times for those meetings. We'll be sending out calendar invites to the metrics work party for those sessions. And if you're in Marrakech, you're welcome to join.

DUANE WESSELS: Yeah. I think that one of our goals from the RSSAC group was – some RSSAC members are very eager to talk about actual thresholds. We haven't talked a lot about that ourselves, and I feel like we're not entirely there yet. But it's possible that during Marrakech, we may have [inaudible]. Also, I'm expecting that we'll have some of the work from our [inaudible] who is taking the RIPE Atlas measurements and trying to apply these metrics to those measurements and see what it looks like. So we may have some [inaudible] to look at. But that's [inaudible].

Warren, can you mute?

WARREN KUMARI: Apologies.

DUANE WESSELS: So, anything else before we adjourn, Steve, Ozan or Russ?

RUSS MUNDY: Nothing from me. Warren does have his hand up.

DUANE WESSELS: I think his hand has been up for like half an hour.

WARREN KUMARI: Yeah, it was getting tired.

DUANE WESSELS: Well, thank you, everyone, for joining the call today. Appreciate the discussion, and we'll see you on the mailing list, and maybe even in Marrakech.

UNIDENTIFIED MALE: Thank you, Duane.

UNIDENTIFIED MALE: Thanks, Duane.

---

UNIDENTIFIED MALE:           Bye.

**[END OF TRANSCRIPTION]**