

---

ICANN75 | AGM – RSSAC Caucus Work Session (1 of 2)  
Sunday, September 18, 2022 – 10:30 to 12:00 KUL

OZAN SAHIN:

Hello and welcome to the RSSAC Caucus Work Session 1. My name is Ozan and I will be the remote participation manager for the session. Please note that the session is being recorded and is governed by the ICANN-expected standards of behavior.

During the session, questions or comments submitted in chat would only be read aloud if put in the proper form as noted in 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 comments verbally, please raise your hand.

If 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.

---

***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.***

---

To ensure transparency of participation in ICANN's multi-stakeholder model, we ask that you sign in to Zoom sessions using your full name. For example, a first name and last name or surname. You may be removed from the session if you do not sign in using your full name. With that, I will hand the floor over to Ray.

RAY BELLIS: Thank you, Ozan. I believe you said that you were going to run a roll call for the session first. Can you hear me, Ozan?

OZAN SAHIN: We can go around the table and ask participants to introduce themselves. So if we could start from the left-hand side, could you please go ahead and introduce yourself?

ABDULKARIM OLOYEDE: Good morning. My name is Abdulkarim Oloyede. And that's me. Thank you.

KARL REUSS: Yeah. Karl Reuss, University of Maryland, Root Server Operator.

BRETT CARR: Good morning. Brett Carr, Nominet.

---

DUANE WESSELS: Duane Wessels from Verisign.

JEFF OSBORN: Jeff Osborn, ISC.

KEN RENARD: Ken Renard, Army Research Lab.

ANDREW McCONACHIE: Andrew McConachie, ICANN Policy Support.

LARS-JOHAN LIMAN: Lars-Johan Liman, Netnod.

ANUPAM AGRAWAL: Good morning, everyone. Anupam Agrawal, Caucus Member.  
Thank you.

OZAN SAHIN: Is there anyone else in the room who would like to introduce themselves? So from support staff, we have also myself, Ozan Sahin. And joining online today, I see we have Baojun Liu, Brett Carr, David Lawrence, John Augenstein, I guess Karl Reuss, Kazunori Fujiwara, Lars-Johan Liman, Paul Hoffman, Peter DeVries, Robert Story, Russ Mundy, Ryan Stephenson, Shinta

---

Sato, and Yoshitaka Aharen. And Duane is in the Zoom, but also in the room. So yes, these are the Caucus member participating today. Back to you, Ray.

RAY BELLIS:

Okay. Thank you, Ozan. Okay, so that's Agenda Item 1, welcome and roll call out the way. So I'll switch on to Agenda Item 2, to discuss a regular meeting time and meeting time cadence. Well, let's get the cadence out of the way first. I listened to the last meeting where we didn't actually get around to discuss this. So I think the options generally are going for a two-week cadence or a four-week cadence. My own preference, I think, to make this timely would be to go for a two-week cadence. Is there anybody who feels strongly either way on two-week or four weeks? So sort of raise your hand. For those in the room, let me know. Okay, Duane, go ahead, please.

DUANE WESSELS:

I think two weeks is good.

RAY BELLIS:

Okay. Absent any other feedback, it looks like we're going to go ahead with a two-week cadence, Ozan. As for meeting time, Ozan, I think it's probably going to be a little difficult to actually pick one

---

here right now. Are we okay to go with a Doodle poll to pick a suitable time? Would you think that would be appropriate?

OZAN SAHIN: So we can definitely circulate a Doodle poll to identify a standing call time for a call every two weeks. Ray, I also wanted to ask when the first call would start in this cadence; so would it be—you know, it's generally considered a travel week right after ICANN public meetings. Would you like to start two weeks from now or—

RAY BELLIS: I would suggest two weeks from now. Yes. So that would be the week starting—one second; week starting the Sunday 2nd of October.

OZAN SAHIN: Thank you. We will take an action item to circulate a Doodle poll.

RAY BELLIS: Thank you, Ozan. I have too many windows open here. Okay. Having got that sorted out, let's dive straight away into the discussion of the label-count as a new metric on the system.

I don't recall who actually proposed this and what the rationale for it was. I would note that, well, if going back to version one, the RSSAC [inaudible] does say that the system is not necessarily

---

used to generate pure research questions. Ken, I see you have your hand up. Please go ahead.

**KEN RENARD:** Thanks, Ray. I think to myself and Andrew, we were kind of initially talking about this. So the idea behind label-count was to look at the adoption of QNAME minimization. I think for discussion here, maybe we should assume for the discussion that the metric is optional. And if we come to an agreement, even if not all operators collect this metric, if it is written down in this document, such that at least the operators that do collect it, choose to collect it, we'll at least be able to have a common format and meaning for what that might be. Does that sound fair?

**RAY BELLIS:** Yep. That does sound fair. Paul, you have your hand up. Paul Hoffman, go ahead, please.

**PAUL HOFFMAN:** So this is Paul Hoffman for the record. Two things on what Ken just said. One is that we still, in Item 5, and I think this will be an ongoing discussion for at least another month or two, what does it mean for a battery to be optional? And so I don't think that that would differentiate this one so much.

I like the idea of having a common format. I'm not convinced that it should be part of 002 as compared to a common format that might be collected otherwise because it really does stand out from 002. It's not an operational measurement. That is, 002 started originally because SSAC said we are concerned about adding too many TLDs to the root zone might cause a noticeable negative effect on the root server system. So let's measure everything we can that might reflect that.

And so there were some initial measurements; others were added. But as far as anyone has said so far, this is not a metric that if the number went up or down would mean anything to any of the operators, that is, it would not mean it operationally. Now, it might, and if it that's true, if some operator says, "Ah, we discovered that in one of the—or the resolver that we're using—I'm sorry, the authoritative server we're using, or one of the authoritative servers we're using, when we ran some tests and did and compared one-label queries to let's say a four-label query or a three-label query, we could see some difference," then sure, put it in. But short of that, I would say it's not necessarily something we want.

David Lawrence asked in the chat, what are the arguments against including it? And there is not a technical argument, it's just another measurement. It puts a little bit more stress on the root server operators as does everything in RSSAC002. It's simply

---

that in my mind, it's not in the same style as what's in RSSAC002 now. And there are measurements, there are other measurements that we might even want general authoritative servers to be measuring that would also not be in the 002 style. And we can easily set up a way for those to be collected and to have them in a standard format. So I am not strongly but leaning against the inclusion of this simply because it's not really part of how 002 is currently. Thank you.

RAY BELLIS: Thank you, Paul. Duane, is that an old hand, or is it back up again?

DUANE WESELS: That's new.

RAY BELLIS: Okay. Go ahead.

DUANE WESELS: So I agree a lot with what Paul just said. I also feel like sort of the connection between label-count and QNAME minimization is a little bit tenuous and not all that well-proven. So if there were to be an addition to 002 for this metric, I would want to see some language in the document sort of saying that don't draw too many conclusions from this because I think just from the way I



---

look at the data, there's a lot of junk in there that's clearly not due to QNAME minimization but could be confused as QNAME minimization. So I would want to see a lot of caveats in the text if we do add this. Thanks.

RAY BELLIS: Okay. Thank you, Duane. Ken, you are up.

KEN RENARD: Thanks. Yeah. So it's not an exact measurement of QNAME minimization, but the long-term trends, if we see fewer and fewer labels of two plus, that's certainly an indication. I definitely agree it's not a part of the kind of current flavor of 002, its initial design there. But I'd make the argument that it somewhat reflects the health of the overall DNS and then the privacy of that. That said, I'm perfectly fine with it not being in there. It's just a possible inclusion that might be interesting again for the sort of overall health of the DNS.

RAY BELLIS: Thanks, Ken. Russ Mundy, go ahead, please.

RUSS MUNDY: Yeah. Thank you very much, Ray. I'm inclined towards including it. And the reason for that I will admit is somewhat back to the

---

point of, I was one of the SSAC people that was pushing for having measurements of the nature that you could, in fact, see if changes or migration of how protocols were being used, how DNS was being used and have at least a consistent way of seeing if, or if not, these changes were having an impact on the root servers.

And so I will also admit, being from a research background, I have an inclination to think having more measurements available is better than not having them available. So even though I've been very involved in operational stuff, I can understand the arguments against it, too. But I think it would be good to try to get something crafted to include.

RAY BELLIS:

Okay. Thank you, Russ. Ryan, go ahead, please.

RYAN STEPHENSON:

Hi, this is Ryan Stephenson, Defense Information Systems Agency, DOD Nic. Sorry, long name. But I'm kind of in agreement with Paul Hoffman and Duane Wessels. And one thing is I would like to see that if it is an optional metric that any optional metrics would be kind of prefaced by a couple statements saying that, hey, not every RSO is a research RSO because, again, we have diverse RSOs. For example, my RSO is more on the site of an ISP as compared to ARL's RSO, which is more kind of like a research-based RSO.

---

But the other thing I would like to kind of point out for this is the fact that with these "optional metrics" that do not necessarily show the operational status of the root server operator, if an RSO was not to publish those, that shows no level of engagement or disengagement of the RSO. It's just that the RSO decides, hey, we just don't want to publish those. That's about it. It doesn't show any kind of semblance that the RSO is not engaged with the rest of the community. And that was it. Thank you.

RAY BELLIS:

All right. Thank you very much. Wes, go ahead, please.

WES HARDAKER:

Thank you. This is Wes from USC ISI. I had a master's student a while ago, three or four years ago, look into label-counts and look into, we actually broke down all of the requests that were coming into DITL data based on the number of label counts. Then we sort of analyzed each one and figured out why.

And my one conclusion from that work is there's so much weird stuff that you have to be careful that the error and the noise is rather overwhelming. And so I do worry that it will obscure the real signal that we're trying to measure or that will make conclusions from a label-count, even longitudinal trend that turn out to be wrong because it has nothing to do with QNAME minimization. It has to do with some other broken firmware that

---

suddenly spouts a lot of data. I mean we were analyzing label-counts up to 50, 60, 70 long because of the weird stuff that exists out there that does leak, especially up to the root.

So you do need to be careful of that. I don't know that I have an opinion yet one way or the other about whether it should or shouldn't be included. We will support adding it. I'm sure, if 002 changes to do it, but I do question the ability to draw solid well-founded conclusions from this particular step.

RAY BELLIS:

Yeah. Thank you, Wes. That's useful information. Just before we carry on with any other questions, I should point out in the text in RSSAC 002. Actually, it's Section 2, the scope of measurements, but the text does say the RSSAC recognizes that measurements of some values are out of scope. Specifically, the goal of this document is not to answer a wider set of research questions, although some of the current measurements may be used in research, time scale, and [inaudible] measurements through RSSAC [inaudible] a poor match to support our research in general. Additionally, more suitable alternatives to support DNS research exist.

So yes, in my own view of this particular metric, I think it is more interesting in the very long term rather than on a day-to-day basis. And if you were to look at QNAME minimization, it could be

---

probably more readily done by looking in DITL data sets than really looking at the distribution as reported by the RSOs. So, Fred, go ahead, please.

FRED BAKER: I just took a look at RFC 9156, which defines QNAME minimization. And there are 31 references to the character strings that are differentiated in that and called labels. So I think some of the comments that have been made suggest that they aren't. And I think not at all true, not at all my observation.

RAY BELLIS: Thank you, Fred. Paul, go ahead, please.

PAUL HOFFMAN: Thank you. So this Paul Hoffman, again. A few things: Russ Mundy mentioned earlier that this is a could be used to measure the health of the DNS. As one of the co-authors of the document, I assure you and as somebody who was active in the working group because of that, the fact that you get some privacy—that you, as a resolver operator, get some privacy advantages by QNAME minimization, that is not related to the health of the DNS.

It improves your privacy. So more privacy is good. In fact, I think we're—even the next topic in on the agenda here is related to that—but saying that an operator or a group of operators who are

---

getting better privacy is improving the health of the DNS, I think, is a stretch that we certainly didn't make in the IETF. Maybe we can make that here in RSSAC. But it's something that we would have to be explicit about and to say why we were doing it.

And to that point, as Duane mentioned, and a few other people have mentioned, that this is a like there may be a relationship between the number of the label-counts and the deployment of QNAME minimization by resolvers. And as Wes pointed out, the actual research that has been done shows that that's difficult to do.

One of the things that would absolutely be needed in order to say this query was a QNAME minimization, is to look at the actual query itself and see whether it is for something that is in ICANN, the IANA root; that is, single strings that are not in the IANA root are not an indicator of QNAME minimization. And when I did a very informal look at the IMRS data and looked at the single-label queries, they were all over the map. You know, plenty of them were clearly QNAME minimization queries because they were for—you know, they were for NS records for com and such like that.

And a whole lot of them, it wasn't clear if they were garbage or not. So, if we include it, we will not only need to put a caveat saying it's not clear how good this measurement is, but I think

---

we're going to have to be honest and say a better measurement would be to do an analysis and do the names. And I can't imagine that the RSOs want to do that effort on a day-to-day basis for RSSAC002.

Fred is correct that 9156 talks about labels. And that's very explicit. 9156 really went down to the label level so that you knew exactly how long of a name to do. So again, that I think supports the strings being labels and such like that. I just think that this still feels like too far for 002 in that sense. Thank you.

RAY BELLIS: Okay. Go ahead, please, Russ.

RUSS MUNDY: Thank you again, Ray. Just a quick response to Paul's comment. I was not arguing for inclusion of this in 002 from any privacy perspective or sort of security and that point of view. It was really that it is—the QNAME minimization is, in fact, a change in how DNS is, in fact, being used as included in the standards. And we don't really know what the impact is going to be. And I think if there are measurements or there is a way that a measurement can be defined, that was really my thought behind, that it would be useful, hopefully, to have it there. And if it wasn't useful, it could be removed in the future. Thanks.

RAY BELLIS: Thank you, Russ. And, Wes, go ahead, please.

WES HARDAKER: Thanks. You know, I guess one other thing to consider with trying to study longitudinal stuff like this is, so I guess my comments are two-fold. What Paul said is probably the right thing to do is measure exactly what it is we're doing rather than do this proxy measurement of labels counts that may or may not accurately represent the problem that we're actually trying to study.

But with respect to longitudinal measurements in the first place, I would argue that this is one that we probably don't need more than one data point a year that could be extracted from better analysis in DITL, right? So and not only that we have the back data available to us through DNS OARC already. And with this, with a more frequent analysis, what are we actually going to do from it, right? We're going to draw some conclusions that maybe it's trending in the right direction. And I don't think we need it every six months or one year is probably sufficient.

RAY BELLIS: Yeah. Thank you, Wes. That actually summarizes my own position pretty much exactly as well. I think the histogram of the distribution is far less meaningful if you don't also know which of



---

those single labels were actually for TLDs and which were not. And trying to put that data into the system to my mind is certainly a very long stretch towards what RSSAC002 was supposed to embody.

Okay, well, I'm not sure how we proceed with the consensus call on this one as such because at this point we've only really got one or two saying they'd like to see it but not necessarily say they really, really, really want to have it. It seems to me there's quite a few saying that there are definitely reservations against including this. From my own position running F root, if this was to appear, I would not expect to be able to collect it.

And I'm also concerned that over the years to come when we start talking about the next question, about DoH and DoT, that we simply won't be able to get that information using DNS packets over the wire because more and more of that information will be encrypted on the wire. So unless you've actually got the instrumentation in the DNS server, you simply may not be able to actually get this information anyway. Liman, go ahead, please.

LARS-JOHAN LIMAN:

Lars Liman from Netnod here. I just want to say that I agree with Wes and yourself. I think I would like, if I were to implement this, I would like an underpinning statement saying that this needs to be used for the following purpose, and that purpose should

---

include expected decisions that are to be made based on this data. So if this data is provided to someone, someone will use it in a certain way, and that will lead to decisions going in one way or another.

Right now, I can't really see that, but if someone can provide that information, I might be convinced to change my standing. But I also think that we have enough data points from the DITL collections. And as Wes said, we probably need to do a lot deeper analysis on this data than just count the labels because there's too much crap in there. Thanks.

RAY BELLIS: Okay, Liman. Oh, Paul, you've got your hand up again I see. Go ahead.

PAUL HOFFMAN: So this is Paul Hoffman again. So Russ put in the chat that it seems that there's at least a weak consensus to not include this. That doesn't mean that the counting should be dead. I think that even if it's not in 002, since there is some interest in the question: can the root server operators see how QNAME minimization, the use of it by resolvers, might be increasing, I think that a new work party that is specifically focused on what sort of analysis from DITL can we do that would be useful, that would generate useful information about the root service is a reasonable thing.

---

So far all the DITL research is being done ad hoc. Sometimes you see a report on it at a DNS OARC meeting, etc. It might be a reasonable thing. And, again, this is not for this work party, but if we are not adopting this, it might be a reasonable thing for the Caucus to take up, is a work party, not to come out with a report, but to come out with a list of interesting research that can be done with DITL, especially that can be done with older DITL and then we'll figure out how to do that and we have plenty of time until the next one and such like that.

I would be interested in that from the research angle, I know there's a bunch of other researchy people here, and maybe that research is in fact pushed by the Caucus or at least suggested by the Caucus as something that would be of interest. So if we don't adopt it here, I think we still have the opportunity as the Caucus itself, not just this work party, to bring this up and maybe find other interesting questions as well. Thanks.

RAY BELLIS:

Thank you, Paul. Do any of the RSSAC members, I guess potentially Fred, in particular, as chair of RSSAC have a view as to whether such a work party is feasible and within scope of RSSAC? No comments?

FRED BAKER:

Okay. So is it feasible?

RAY BELLIS: Oh, go ahead, Fred.

FRED BAKER: We could schedule meetings related to that. I'm not sure exactly where your question is going.

RAY BELLIS: Okay. All right. On that basis, I think I'm going to have to say that we do have a weak consensus not to include this in RSSAC002. Let's get that recorded for the record. Right. Let's put this in the agenda. Okay.

So moving then now on Item 4, which is discussion of possible additions of metrics to DoH, DoT, and DoQ—that being DNS server HTTPS, DNS over TLS, and DNS over QUIC. Do we have any proponents or otherwise for this possibility? Go ahead, please, Duane.

DUANE WESSELS: Thanks, Ray. This is Duane Wessels. Can you clarify—because the way it's worded here is a little bit different than I was expecting. What I was expecting was that for—we already have some metrics where we count by transports, and I think the suggestion is to basically add new counters for these new transports, which is

---

maybe a little bit different than what is being suggested here. Or is that the same thing?

RAY BELLIS:

I would expect that these would appear in the metrics that are currently split out by layer three and layer four protocols because we already have counts based on all four combinations of IPv4 versus v6 and UDP versus TCP. But we also, you have—in fact, from memory, it's actually [on the bucket count.] Sorry, no, that's not true, is it? I believe we actually have both [summary] totals and packet size-based totals for these at the moment. But, yes, I would imagine that if these would go ahead, they would get retrofitted into there. Yep.

DUANE WESSELS:

So it sounds like it's adding new keys to existing metrics rather than adding new metrics.

RAY BELLIS:

That would seem the logical implication of that. But that's not necessarily the case.

DUANE WESSELS:

Okay.



---

RAY BELLIS: Thank you, Paul. Go ahead, please, Wes.

DUANE WESSELS: Thanks. You know, I think you bring up an interesting point, Paul, that there is not a standard. I question whether that means we shouldn't measure it because there's not a standard when we know there are cases of people starting to do it. There are companies that are starting to probe.

You know, a lot of this came from the presentation I gave at the last DNS OARC where we implemented TLS to our root server as an experiment. I have architectural plans on how to do that safely and possibly open it to the public in the sometime near future as an experiment. And the question came around when I was looking into this of, well, if I do turn this on, how do we report things that are in 002 that no longer match? So the transport strings, UDP and TCP are there in the past, and we don't have these three, as I think has been mentioned.

It would be interesting to note I think one of the things I suggested is we take 002 to make it easier to add more prefixes in the future. I don't know what that would be or if the operator in question that accepts a new transport would accept it. I don't know.

It seems to me like these would be easy to add, and though there is not a standard for how to do it, that does not mean that it's not being done, right? The IETF is often very late in defining a

---

standard for something that already has running and sometimes even deployed code. The fact that we want running code in the IETF before we publish a standard sort of indicates that implementation occurs first. And that's not, of course, how the IETF always works, but nonetheless, that's one of our goals.

I will note that these are not the only things that I identified as a problem. So the other things that I don't think are on the list yet include traffic volume and how you measure that when the size of the packet is inside an encrypted tunnel. Unique sources, if I open a TLS connection to a root, but I don't actually send a query, do you include that source, right? Just open a connection and close it. That's true for TCP too. I actually don't know how we deal with that. It's sort of a new problem I hadn't thought about.

Our code volume, that our code is now encrypted. So it's no longer as easily visible. I mean, there's a slew of problems that kind of amount, but this is the biggest one in my mind, of how do we count those requests? And regardless of whether there's a standard that points to it, if it's implemented, turned on, and in use, I'm not sure that we should say there must be an RFC before we start counting.

RAY BELLIS:

Thank you, Wes. Go ahead please, Fred.



---

FRED BAKER: I find myself thinking about the QNAME minimization discussion that we just had. And it seems to me that the same arguments would apply to DoH, DoT, and DoQ. And if you don't want to do it for one, you shouldn't do it for any of them, for the same reasons.

RAY BELLIS: I see where you're coming from, Fred, but I think then that might send us into the sort of pit that Paul described in the last meeting, where if we say we're not doing this, then we can start to question why we're doing any of them. Yeah. Okay. Paul, go ahead.

PAUL HOFFMAN: So I would disagree with what Fred just said on a very fundamental level, which is that if there's one or more root server operators who are offering any of these three transports, it's going to have an operational effect on them. And it's going to have a much more significant operational effect than, for example, if the number of TCP—just plain old port 53 TCP queries quadrupled tomorrow. That would be somewhat devastating to many people. But if you also included then TLS or QUIC handshakes, that would be noted significantly.

So if these things are going to be offered, they will have an operational consideration. And so I will now not reverse myself, but I will speak against what I said earlier. If these are going to be

---

offered, we should be counting as soon as possible, even if we are counting things that we aren't sure what we're counting.

And I put a note in the chat. Right now we're assuming we know which port numbers are going to be used by DoT and DoQ. That could easily be wrong. We got some significant pushback on saying, oh, we're going to use the obvious port numbers. People are like, "No, you shouldn't use the same port numbers that stub resolvers are using to the cursive resolvers. You should be able to differentiate. There are plenty of port numbers available.

So, to me, I am not sure. I'm leaning strongly against it unless there are going to be RSOs who are going to be offering it and possibly having operational considerations. And then we can add it. But I think we're going to have to grind hard on the wording to say, what are you counting? And as Wes pointed out earlier, the what are you counting gets really difficult, even for like questions of, did a query get sent? Did you respond to the query? You know, because a TLS handshake or even a QUIC handshake can be aborted for a zillion reasons beyond what one can abort a TCP handshake for. Thank you.

RAY BELLIS:

Okay. Thank you, Paul. Well, Ken, go ahead, please.

---

**KEN RENARD:** One thing that could come out of this is at least in the document updating the text to clarify that DoH and DoT should or should not be considered TCP measurements or as in the TCP count. Same with the DoQ and UDP. So there is no ambiguity of whether—if RSO is supporting these, that these counts go into the TCP bin or they should—they go into whatever at least defining what the behavior is.

**RAY BELLIS:** Yes. Thank you, Ken. That does seem like a very suitable suggestion that we should only clarify that the TCP counts are not inclusive of any DoT-type queries or anything else that happens to be using a TCP layer three underneath. Wes, go ahead, please.

**KEN RENARD:** Yeah, the double counting. I hadn't thought about that problem. That's a good one, Ken. Because DoQ is the same thing. You count that in UDP or not. But actually, the reason I raised my hand though, is to remind, I guess, everybody that even if authoritative requests to servers was defined by the IETF, whether or not it should be deployed at the root is, A, subject to each RSO, but B, that actually would probably come out in either a different RFC or more likely an RSSAC document saying the root server system should now support new protocols or new transports. I think the implementation of TCP predates sort of our policy stream. So we

---

didn't do that back then, but there was a long rollout for AnyCast. There was a long rollout for DNSSEC. And just because the standard exists doesn't mean it's going to be instantly available because there's operational considerations that must be taken into account.

WES HARDAKER:

Yeah, the double counting. I hadn't thought about that problem. That's a good one, Ken. Because DoQ is the same thing. You count that in UDP or not? But actually, the reason I raised my hand, though, is to remind, I guess, everybody that even if authoritative requests to servers was defined by the IETF, whether or not it should be deployed at the root is: A, subject to each RSO; but, B, that actually would probably come out in either a different RFC or more likely an RSSAC document saying the root server system should now support new protocols or new transports. I think the implementation of TCP predates sort of our policy stream.

So we didn't do that back then, but there was a long rollout for any cast. There was a long rollout for DNS sec. And just because the standard exists doesn't mean it's going to be instantly available because there's operational considerations that must be taken into account.

---

RAY BELLIS:

Yes. Thank you, Wes. I would also like to back that position as well, that I think, or at least I think I'm agreeing with you that until such time as the RSOs collectively agree that this is something that we're offering, I personally feel it would be premature to start specifying this. But I still agree with Ken's earlier statement that we should clarify that TCP really means TCP Port 53 and not anything else that happens to be wrapped up in a TCP session. Peter, go ahead, please.

PETER DEVRIES:

Yeah. Hi. I think I agree that this is different than the other subject—label-count—in that this can affect performance and this can affect—this is lessons learned that RSOs can directly implement. So I think that metrics like that should be considered differently. I think on the subject of these specific ones, maybe it's enough to put something more generic that says if you're implementing some new transport or something like this, you are expected to voluntarily publish some metrics on that that you decide and that those will be clarified in a later document, you know? So something to say, to capture that data because I think it's valuable for others that might cross that path in the future. But something that just doesn't tie us to these specifics. Okay. Thanks.



---

**DUANE WESSELS:** Thanks, Ray. This is Duane. So I wanted to bring up something which maybe needs to be a separate discussion point. But as Wes was just talking about adding these and making these optional. And, in fact, the current 002 document still says that even TCP itself is optional because reassembly is hard, basically. So the bar is already pretty low for RSSAC002. An RSO is not required to publish any TCP metrics. Now, I think we all probably do, but I wonder maybe if we want to revisit this aspect of 002 and make TCP no longer optional.

**RAY BELLIS:** Okay. It seems like it's too early to make a call on this, but I think Duane's possibly giving us a perfect segue into the next agenda item, which is discuss the optionality of metrics in RSSAC002, and specifically, what does it mean for a metric to be optional in the first place? That's the 5a agenda item. As others have noted previously, actually nothing in RSSAC002 is completely mandatory for any RSO to implement. And in so far as we've generally—as RSOs we've committed to do so. You know, we will make best efforts to produce the best way we can. But nothing is actually strictly mandatory for any of us to produce. So does anybody have any thoughts on this optionality question yet? Paul, go ahead, please.

---

PAUL HOFFMAN:

So this is Paul Hoffman. We are having this very discussion about RSSAC001V2, and I would like to see the two documents aligned however it goes, and I don't think we'll be in a deadlock. I had proposed some new wording for 001V2 that Wes made a good correction on and such, but basically, the idea is at least where we're started and I guess we'll have more discussion on that later this week is using the word expectation throughout 001 instead of expectation, requirement, recommendations and such like that. And then to have a clear definition of what an expectation is.

I think that it is perfectly reasonable for 002 to say there is an expectation that the root server operators will do all of these where an expectation means “And we know that some of them won't be.” And I believe it will probably end up in a similar situation for 001V2. I don't think that we have to have different kinds of exclusions, but we will clearly have different exclusions because stuff from 002 isn't in 001. But I would like to make this document match whatever comes out with 001. If they aren't matched, a non-RSO who's reading them, especially somebody who's looking to see what are the RSOs doing wrong. Or I'm sorry, what is an individual RSO doing wrong, is going to have a real hard time like justifying, “Oh, look, they updated these two documents at the same time and they don't mean the same thing.” I would really like to see them mean the same thing.



RAY BELLIS:

Thanks, Paul. And yes, I've noted your comments from the last session, which I listened to earlier. Yes, it's a tricky one. It struck me looking at the documents and reviewing the last meeting that possibly what we need, and this also actually touches on 5b, is a standard way or a standard place at which each RSO may, not necessarily required to, publish their deviations from the expectations of RSSAC 002.

I know there were some questions about the unique sources and F root. But I'll go with that again just so it's actually on the record. At F root, we publish the unique sources as seen on our own instances, but we do not collect unique source counts from the Cloudflare instances operated under contract by them. And in part, that's because it's too onerous for them to collect the individual v6 unique sources or unique prefixes at least and then perform the union because this was discussed actually a couple of versions ago, that at first approximation, you might argue that the set of sources served by Cloudflare plus sources served by [inaudible] F root, the sum of two sets is to a first approximation the same as the sum of the unit of those two sets. But it's not necessarily the case.

---

So, to do it properly would require that they actually send us all the prefixes and then [inaudible] prefixes. We then combine the two and take account of those, that combined set. But we simply decided it's just not practical. But I'll be quite happy to have a file somewhere on RSSAC 002 repository that says, "This is the data we have. This is the data we don't have." Any thoughts on that from anybody? Wes, go ahead.

WES HARDAKER:

Well, so, let me bring up question C, which is missing from the list is, if you are not able to count something in the way that the text is defined, should you report it all, right? So, that's sort of the inverse of optional. In other words, if you're unable to do it according to how it's specified, then you probably shouldn't report it as opposed to under-report it, which actually even goes back to the TLS and TCP thing, right, which is the if it comes in over TLS, and I can't count it properly, then I shouldn't report it at all.

RAY BELLIS:

That's an interesting point. Duane, go ahead.

DUANE WESSELS:

Thanks, Ray. I think your suggestion is good. I think it would be helpful to have some way to know that when you're looking at the

---

data, if the omission is intentional or not. My preference would be to put something in the data files themselves that says, “We opt not to publish this metric for the following reasons,” or whatever, or, “We aren’t able to.”

Additionally, to the broader point about this topic, discussing optionality of metrics, I think we might be confusing the idea of optionally publishing metrics versus publishing nonstandard metrics, which is kind of the discussion we had earlier about maybe the label count. So, somebody might want to publish something that’s not documented as a standard metric, which they should probably be able to do.

And then, there are cases where metrics or certain parts of metrics may not be included because, well, for various reasons. So, thanks.

RAY BELLIS:

Yes. Thank you, Duane. Now, I do recall that we have text in Version 4 of how to... Oh, wait, Section 7 identifies specific metrics. So, this is how an RSO would publish a metric that is specific to them, and not part of the core set. Let’s see if there’s any ambiguity on this one.

UNIDENTIFIED MALE:

Yes, Section 7.

RAY BELLIS:

Yes. Yes, what we don't ... It does seem to me that Section 7 could be a little more explicit about what the resulting file names would look like for those to identify specific metrics. I'm not sure we can necessarily assume it from the text that's there.

Does anybody else feel that text could do with a little bit of beefing up there to talk about what the expected path would be? Actually, it's possibly covered by Section 6.7, which is the immediate preceding section. So, it talks about a... Oh, metrics [inaudible] and then a short service name. And a short service is... Oh, okay. Just [x-root.] That's fine.

And I think 6.7 actually probably covers the case I was concerned about, not knowing what the actual path would be. Excuse me. I lost my train of thought on that one now. It's 4:30 in the morning here. At least haven't got jetlag to contend with. Let's get back to the agenda. Yes, Paul, I've certainly seen all of the feedback here in our last session about basically trying to sync in with RSSAC 001. I haven't actually yet been attending that group. I'm going to actually have to start doing so.

Yes, okay. Yes, Duane, you raised interesting points about whether we should actually try to have something in the data itself that says whether or not—describes any nonconformity. I'm slightly concerned that if we try to invent a schema for having

---

something that can be read by a program, that we may actually struggle there. I know we have the case that's been discussed at the last meeting about, well, what happens if your v4 or v6 addresses have been obfuscated before they're fed into the data capture pipeline.

I don't know whether any RSOs are actually doing that. So, is there any RSO actually here on the call that can say that they're actually processing the unique sources in such a way that they are affected by obfuscation or in their process and pipeline, or can you say that they're not affected? I can say for my own part on F-root, our data processing is done before anonymization of IP addresses. Duane, go ahead.

**DUANE WESSELS:** Yes, it's the same for Verisign. The metric is not affected by any anonymization for us.

**RAY BELLIS:** Can any RSOs confirm that or not?

**KEN RENARD:** H is the same.

**KARL REUSS:** D is the same as well.

---

RAY BELLIS: Yes. Thank you. That might be something that we can actually survey all the RSOs to actually get confirmation of that to find out whether this is actually an issue at all or if all of the RSOs are currently managing to measure without obfuscates, then, I think the issue that's come up at the last meeting can be deemed a non-issue in terms of whether there's IPv4 addresses that are actually valid that are /24. I'm not aware of anybody that's actually said at any of the meetings that they are not measuring those.

Actually, I'm going to put a question to Andrew McConachie. I know he's done a lot of processing on the RSSAC 002 data. Have you seen anything to suggest just while perusing the data that any RSO is seeing significantly fewer IP addresses than the other RSOs that might be accounted for by obfuscation?

ANDREW MCCONACHIE: You mean for [known]-sources IP4 and IP6?

RAY BELLIS: Yes.

ANDREW MCCONACHIE: I actually hadn't looked at it in a very long time. I could look right now.

---

RAY BELLIS: I guess you're looking at the same link that Ken just posted to the chat. Curiously, there's a very large jump in the number of unique sources. Oh, [inaudible] v4 versus v6. So there was a very large drop in the number of v6 addresses being reported just before the start of this year.

ANDREW MCCONACHIE: So, if you look at queries over sources, it's kind of interesting because the thing you mentioned about F Root or because—

RAY BELLIS: Yes, because we're actually under-reporting the number of unique sources. So, if you're dividing by that, you're getting actually larger numbers than would otherwise be expected.

ANDREW MCCONACHIE: Exactly, yes.

RAY BELLIS: Interesting. And Wes was suggesting that as maybe an argument for simply not reporting the data instead of reporting complete data. Paul, go ahead, please.

---

PAUL HOFFMAN:

So, I think this is an interesting question. I'm not sure if it's one to be answered in 002, or I'm sorry, I'm not sure if it's one to be answered in the meat of 002 but, instead, something in the introductory materials saying, "There are these questions. Some people have obfuscation and such." I think it would be good to note it.

But I think it would be bad to come to conclusions because any conclusion we come to could be obviated by BIND or KnotDNS and such changing the way they do things in six months or whatever. And then, we would be getting a mixture. So, I would prefer that if we deal with this, and I think it's reasonable to deal with, we deal with it in text, not in the collection part per se.

RAY BELLIS:

Okay. Thank you, Paul. I think it would be prudent to actually survey all 12 RSOs and get explicit confirmation of whether obfuscation is currently an issue, or whether they are doing anything—aggregating the IP addresses in any form other than v6 where we already specify that aggregation is per /64. I'm going to bounce back to Duane.

You were talking about having some means of specifying nonconformity within the data files. Maybe I don't really have any chance to think about that first since you mentioned it a few minutes ago. Do you think such a task is actually feasible?



---

Because I do have concerns that it would take longer to make an electronically readable nonconformity than it would be to simply have a text file that says, “This is what we do.”

**DUANE WESSELS:** I guess the reason I suggested having it in the data file is so that temporally, it’s bound together, right? So, if you had an external text file, you wouldn’t necessarily know, do the contents of this text file apply to the data at the time period X, Y, or Z? If I was to make a proposal for doing it in the [inaudible] file itself, I would suggest just to have a notes field or something like that, a free-form notes field, where they are still expected to explain any differences from the documented expectations.

**RAY BELLIS:** Okay. So, you’re not necessarily expecting any processing to be changed as a result of that, or at least the nonconformance would be documented there in the file so that somebody could read it if they needed or wanted to.

**DUANE WESSELS:** Yes, exactly.

---

RAY BELLIS: Okay. Thank you. Yes, I think that’s definitely something we could attempt to do. Russ, go ahead, please.

RUSS MUNDY: Yes. Thanks, Ray. Yes, I was going to suggest something similar to what Duane is suggesting, I think. And that is that the indication in the data file itself could be extraordinarily simple where the RSO, if you believed it was in conformance with the description of the what’s being collected, [inaudible].

And then, there would be maybe another field that could point to the URL where the text file describing exactly what was being done, if it wasn’t in conformance. So, it seems like the data content itself could be kept very simple and pointed to another related URL that gave explanation.

RAY BELLIS: Yes. Okay. Thank you. Okay. What's Peter’s comment? “I think prefer don’t report over reporting [inaccuracies or comments, very quickly get to a point where there's zero data being reported.]” Interesting point. I’d like to think that most of our data is very accurate, but I know the unique sources H and F root is not accurate. Yes. Everything else should be. Russ, is that a new hand? Oh, okay. Robert, go ahead.

---

**ROBERT STORY:** Just this discussion about freeform text fields and I guess describing variations had me harkening back to my days in SNMP and mid-modules that had defined agent capabilities and then compliance where we could say, “This is what we expect this data is maybe with some variations in the strictness or how it’s collected.” And then, the data file could say what level of the document those definitions are coming from, and then what level of compliance, or which items you are compliant with. I’m just throwing that out there.

**RAY BELLIS:** Okay. Yes, thank you, Robert. I’m reminded constantly that any protocol that says it’s simple and lightweight inevitably turns out to be neither. I think we’ve probably got enough there that we might in a position of where we can start forming some text around this a little bit. We’re not going to do that on the call. I know Andrew told me earlier we haven’t yet got to the point of there actually being a draft of the 05, but there will be, I believe, very shortly. So, that’s I think something that he and I could start drafting with some text for the next meeting. Andrew, go ahead, please.

**ANDREW MCCONACHIE:** Yes. I think you just said it. So, I should just start V5. I’ll just copy V4 in a Google Doc. We can start editing it.

---

RAY BELLIS: Yes, I think that's [inaudible]. And I think, well, whether [inaudible] transcript or whether ICANN's been taking notes, I think we've got a decent list of things we can actually start to put into the document. Does anyone else have any feedback on Section 5 before we go ahead onto any other business?

Okay. I think we're going to call Section 5 concluded for now then. And Andrew and I put some text in time for the next meeting. And let's go to six. Any other business. Does anyone else have anything else to raise on this?

I do have one myself, which is looking at the statement of work, I know this was discussed briefly in the last session, do we feel we actually need to do a section-by-section audit of the existing V4 document just to actually say, are we happy that we've reviewed all these sections and are confident that they are still relevant and appropriately specified? Go ahead, Paul.

PAUL HOFFMAN: I would love to see that audit done by somebody who is not an RSO. I would love to see such an audit done by one or more people who are interested in the root server from the outside. I think that we, the sort of larger we, have gotten used to everything there, especially the RSOs, who are producing the

---

reports. So, I would like to leave that in. As long as we can get... If we ourselves are going to review it, no, I don't...

Yes, we can go through it, and I don't think we will do a good job, per se. But if there are work party members who are not RSOs who are here because they're interested in the RSS, and especially in measurements, having them, having at least one of them, but maybe a group of them, doing it individually would be really useful to the document itself. I strongly suspect that there will be at least one finding where the rest of us are going to go, "Oh, my God, I can't believe we have had that in the last three versions." So, yes, please, let's leave this in. if we can get somebody interested in it who is not somebody who's already implemented. Thanks.

RAY BELLIS: Thanks, Paul. Interesting idea.

FRED BAKER: Yes. Listening to your description of who the somebody might be, that sounded like a definition of the caucus. Do you have anything more specific that you would like to say in that regard?

PAUL HOFFMAN: Yes. I meant individuals in the work party or in the caucus, not the caucus as a whole, because the caucus as a whole has many

---

people who work for RSOs and such. Specifically, I would like to see the review done by somebody who is not an RSO. And we have plenty of those. There are some on the call today. There are others who have participated in the caucus in the past. I think if we state our request that way, we might get volunteers.

RAY BELLIS:

Yes. Thank you, Paul. The statement of work, item one, does specifically say review existing RSSAC 002 measurements to determine whether they require updating to accommodate changes in DNS technologies. And whilst there's been a little bit of discussion on the caucus list and on the last call about whether there is anything that does need to be changed, I do think a more structured review is certainly appropriate to satisfy [inaudible].

I don't think there's any reason why that shouldn't also be done by people who are RSOs. But if there are potential volunteers in the group listening now, maybe that's something we can actually put to the caucus mailing list. If there aren't [inaudible] volunteers here, I think we should go ahead and do that. So, if we do have anybody here that would like to take that on, please let me know and say so now, if possible, please.

Okay. I'm not seeing hands go up. So, that sounds like something [inaudible] we might need to take to the list probably early next week once the ICANN meeting [inaudible].

---

Okay. I noticed a query in the chat from Suhayb about where's the script that generates YAML files in the root server. And as Paul has said, well, it's actually somewhat RSO-specific. So, I'll come back to you in a moment, Anupam. Every RSO will have different means of capturing data. I speak for F root. We are using DNS Cap, but only for unique sources. We use the [RSSM plugin.] We used to use proprietary code that I'd written. But we've recently replaced that with DNS Cap [and RSSM plugin.]

We're not using DNS Cap for the other metrics, which we're actually collecting directly from BIND. So, we have scripts that pull the BIND HTML [inaudible] and collects the data. We have complicated scripts that get that data on each node and then other scripts that aggregate the data from all the nodes into one single file. And you're not likely to find those scripts published anywhere because they're just simply internal parts of our workings. Anupam, go ahead, please.

ANUPAM AGRAWAL:

Yes. Thank you. If there's sufficient interest, I would like to volunteer for looking at the metrics. I'm definitely not from RSO and have been engaged with RSAC 048 revision. Thank you.

RAY BELLIS:

Yes. Thank you, Anupam. That's appreciated. Do you think you might have time to perform such a review before our next

---

meeting in approximately two weeks' time, or would you need a little bit longer than that?

ANUPAM AGRAWAL: I need to have a little more time. But I can submit an interim report in the next meeting in two weeks. Thank you.

RAY BELLIS: Okay. Thank you very much. That's great. Okay. Do we have any other AOB questions or issues to raise? Okay. In the absence of any other AOB issues, I think we should close the meeting. Ozan, if you want to do the formalities there, please?

OZAN SAHIN: Yes. Thank you, Ray, and thank you, everyone, for participating today. We are now adjourned. Tech colleagues, please stop the recording.

RAY BELLIS: Yes. And we'll get the Doodle poll out as soon as possible to arrange the time and date for the next meeting. Thank you, all.

**[END OF TRANSCRIPTION]**