

---

ICANN74 | Policy Forum – RSSAC Work Session (2 of 2)  
Wednesday, June 15, 2022 – 16:30 to 17:30 AMS

OZAN SAHIN:

Hello, and welcome to RSSAC Work Session 2. My name is Ozan Sahin, and I am the remote participation manager for this session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior.

During this session questions or comments submitted in chat will be read aloud if put in the proper form as noted in the chat. If you would like to speak during the session, please raise your hand in Zoom. When called upon, virtual participants will unmute in Zoom. Onsite participants will use a physical microphone to speak and should leave their Zoom microphone disconnected.

For the benefit of other participants, please state your name for the record and speak at a reasonable pace. You may access all available features for this session in the Zoom toolbar.

We have an overflow room called Kilimanjaro located across this room. If this meeting room reaches its full capacity, ushers will help additional in-room participants to the overflow room. With that, I will hand the floor over to Ken Renard.

---

***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.***

---

KEN RENARD:

Thank you, Ozan. Welcome, everyone, to work session 2. Those online as observers as well, welcome.

Today we're going to talk about two different documents and updates to each. The idea is that we're going to see what we want to do as far as spinning up work parties to update these documents. The documents are RSSAC002 and RSSAC001. The updates to 002 are very minimal, but we also want to open up the floor to other potential items to discuss that are updates to that document.

I guess we can go right into 002 if there are no questions. So the proposal here, the main one of two items, was adding a label-count metric to the RSSAC002 stats. The idea being that we could try to track the adoption of Qname minimization.

I believe we've sort of missed the first wave as far as watching that but can we at least see something measurable. Qname minimization is a very effective privacy measure at specifically the roots. There are other privacy measures as well. So I've been playing around with this myself and if we can go to the next slide I can show you some of the tools that are available to measure this.

So the DSC and PacketQ, those changes have been committed and there are minimum version numbers for this. Personally, we use dnscap and the RSSM module with that. So I've added the

---

label-count metric. I just have not initiated the pull request yet with the OARC folks.

But the first two tools, the DSC and PacketQ, they just give you this label-count, how many queries were received. So it's just a full distribution. The RSSM module will actually include the tail in the last sort of bucket, the highest order bucket. So the input of what that N should be.

In theory, for Qname minimization we need zero labels, one label, and two labels. That's all you need to track Qname minimization. But the actual cost of doing more buckets is extremely minimal.

I think on the next slide I actually have a little bit of data I collected over about a week. I think this was a week or maybe a couple days. So on the left side is just the direct output of a dnscap run. You see that for 16 buckets that last bucket, 15, includes the tail and then aggregated. So kind of standard stuff.

And next slide I have just a little data that might tease at some of its usefulness. So this was a month's worth of data from that dnscap module. Just look at the distribution. So there's quite a bit of just single-label TLD-only queries. Kind of an expected hump at three labels and then it trails off pretty nicely.

But it's nice to see that 36% of queries where the label-count is at least one or exactly one. So we're definitely seeing some Qname minimization out there.

---

Next slide shows a little bit per ASN. And this isn't a part of the proposal for RSSAC002. To me it looked interesting data. The top graph there is all ASNs. This was about a week from most of our instances. So there are a fair number up there, about 1,000 ASNs that were 100% single label. But they're only sending one or two queries in that week.

So if you go down to the bottom graph those are ASNs over that time period that sent over 10 million queries. And there's a couple out there close to 100%. So that's good. So this was done with the PacketQ tool which can look at label-count as well as ASN and the source.

So I think the next slide is just more pretty graphs, kind of what you'd expect. This is just looking at zero, one, and greater than one labels over a month period, no noticeable shift or anything like that. But I think over time, over months, years, we would hopefully see the effect of Qname minimization adoption swapping those top two lines.

And the next slide, I think that goes on to the load-time. Andrew.

ANDREW MCCONACHIE: Yeah, thanks, Ken. Oh, Wes, did you have a question about? Okay, sure. So this is the other thing that would need to be updated in RSSAC002v4 or the new version v5 is there's just some clarifications around the load-time metric. On page 11 of v4 it

---

says, “If the load time metric is unavailable, it should not be listed,” and it's kind of unclear if that means that the file shouldn't be present or if the zone serial numbers just shouldn't be present.

So I think the spec could be a little bit more clear there. I was looking into this. Because a few months ago I had a conversation with Howard Kash and Duane Wessels. Where Howard was updating some of their RSSAC002 data to remove zeros from there. They had reported load-times of zero. And it was unclear how a zero should be interpreted by the API that I had written and the resulting graphs. I mean it's a rounding question ultimately. Like you have a load-time somewhere between zero and one and do you round to zero or round to one?

So what he ended up doing was he changed all of the zeros to ones. So there were no longer any load-times of zero. I think there are still some zeros out there in like older RSSAC002 data. And there's nothing saying you shouldn't ever report a zero. So, yeah, it would be interesting to have a discussion about whether zero is an acceptable value. Yeah, so that's another thing to clarify here. Because right now it just says report an integer. And that's it.

KEN RENARD:

So ultimately the goal here is to create a statement of work and start up a work party. But it's absolutely valid to discuss some of the issues here. So, Wes.

WES HARDAKER:

Thanks. So, yeah, I mean, I think that we should open the work party. I think there's interesting things to do. To some extent, I think – something you said in mail was we're too late to the metrics party. And that's okay. That doesn't mean that these wouldn't be interesting things to watch over the future anyway.

I think Duane did some analysis. It might have been Duane with A and J over a long time. And you can see a beautiful graph of the increase of a single-label TLD to show Qname minimization. So even though it might be too late for the beginning history of that, we can see how long the tail is at least. But we shouldn't necessarily debate the right way to do that now.

There is some other things that I think go into that problem space beyond just label-counts. For example, label-counts within existing TLDs versus label-counts of garbage. If you look at the label-counts of garbage, there was a huge amount of single labels back in the Chrome days, right? So if you don't separate those out you may be tweaking your statistics some.

So I think the workspace is bigger than just the simple statement of do we want to do label-counts because of that all means.

KEN RENARD:

Label-count in X domain versus label-count and no error or label-count and no error versus—

WES HARDAKER: Right, that's the cheap way to do it, yes.

KEN RENARD: Okay, yeah, that's a good point. Any thoughts on... Paul, thanks.

PAUL HOFFMAN: Greetings. So I have one comment, one question. The comment is on collecting the label-counts and – as Wes just said – the analysis. 002 actually doesn't do any analysis. It just collects data. So if we are expecting to add analysis to this, I think then the work party would need to do a lot more work. Because there's analysis on all the other metrics as well. If we're just collecting data that's just fine.

Plenty of that data goes into other analyses, which leads me to my other question. Which is will the work party also be open to other things that are not in this current list of two things? For example, although I don't know whether it has come up as an issue, the unique sources metrics is listed as optional in the current version. And it's not clear what optional means and whether someone's supposed to do it.

So if we find other little bits like that, will that be part of the work party? Or would we need to change the charter or something like that?

---

KEN RENARD: Oh, I think those would absolutely be in scope. To revamp the entire document, whatever we see necessary. I think at this point bringing up things just like that would be to add them to the statement of work to say, yes, work on these specifically and then add others as the work party sees fit.

For example, I just talking earlier about IDNs, universal acceptance. Is that something worthwhile or IDN TLDs. Is that worth collecting? Is that interesting data for the purposes of our RSSAC002? It's interesting from a research perspective. But is it something that RSSAC002 should do? Lyman?

LYMAN CHAPIN: Thank you. I actually have a couple of things. You never paused so I never got a chance to ask. In your previous presentation, the tools you mentioned, are all three of them hosted by OARC: RSSM PacketQ and DS—

KEN RENARD: Yes, they are.

LYMAN CHAPIN: Correct, thanks, and just—



---

KEN RENARD: It's a dnscap that hasn't been pulled yet, right. Hopefully next week.

LYMAN CHAPIN: And you're listing statistics by the number of labels. I actually think it would be interesting to expand that list a bit before you do a tail bucket to 34 labels. Because that's the number of labels for ipv6 reverse lookups. So just a hint for future analysis that maybe that's something you want to keep in mind.

KEN RENARD: One of the reasons I haven't done the pull request was I wanted to know what that end could or should be.

LYMAN CHAPIN: Okay, then I suggest not below 34 but it's a suggestion.

KEN RENARD: Yeah.

LYMAN CHAPIN: And then we come to modifying this document. First I believe that – with Paul's words in my ears – we should not go into analysis in this document. Let's keep it simple. And have a document that only lists what we're supposed to collect so that we can easily map what we do against what the document says. And how that

---

data is used later on, that should not be prescribed in its document. It should be a fairly open field for people to do analysis on these numbers, that they shouldn't be prescribed.

That said, having input from people who are interested in certain metrics to use them for certain things is fine. But I think that we should try to limit ourselves to metrics that we can foresee have an impact on the day-to-day operation of the servers. So metrics that can be measured, doesn't have to be measured day by day in order to keep the operation running.

That can be handled with the DITL collections and things done more seldom that doesn't consume power from the servers as we operate. But if there are things that lead to quick operational decisions, then that's something that we might want to have that. But seeing the impact of IDN is maybe not something that it has a direct operational impact. So I would suggest we leave that out. But if people can motivate why we should have them in, then that's fine. Thanks.

KEN RENARD:

Agreed. And I'd ask the same question about label-count. And it's for the work party. But it's not an operational concern. It's more of a how is Qname minimization doing.

---

LYMAN CHAPIN: Yeah, you have a point there. But again that's for the work parties to discuss.

KEN RENARD: Thanks. Wes.

WES HARDAKER: Thanks, and thank you, Lyman. And we should only consider things that are fairly easy to calculate. We've run into this issue in the past with 002 where people have wanted interesting research data that would be fantastic to study. And I agree with Paul. The analysis shouldn't go into 002. The purpose of 002 is to find metrics that can be analyzed, not are analyzed.

So I guess my recommendation for this would be let's create a work party document that says in the beginning of the work party we will do some brainstorming. But try and keep the number of things defined within some scope. And the two caveats that I think Lyman and I just reiterated, right, which is they have some operational context and are not challenging to produce. I think one of the ones in the past was like per IP address or something like that. And that was challenging to produce.

And then part two would actually be doing it. We can list some examples of things we know now might be interesting to calculate like labels. I'll note that doing internationalization is an interesting problem. It's very easy to calculate if the first four

---

letters are XN--. It is much harder if you have to validate that string to prove that is an international label and not garbage, so we don't.

LYMAN CHAPIN: But you may have to search for that label because it may not be the first label. And then you're into processing again.

WES HARDAKER: Well, you're already counting labels, right? So—

KEN RENARD: All good points, yeah. The IDN, it's a long shot. It's certainly not operational. It would just be interesting, yeah. And not trivial to collect so. Okay. A hand from Paul. Paul, thank you.

PAUL HOFFMAN: So two issues and I don't know whether this is relevant to the charter or not. But recommendation one in RSSAC002 says the RSSAC recommends each RSO implement the measurements outlined in this advisory. So this is still a recommendation, not a requirement. But there is an assumption these days that everyone does 002 because there it is there.

But I don't know whether that you want the Caucus to be evaluating whether to change that into a requirement. I know

---

that it is actually relevant to the next topic which is RSSAC001. But if we also take up the work for RSSAC001v2, which we'll be discussing next, the question of requirement for collecting the measurements for 002 I think is relevant.

And I do believe that it's a bit dicey for exactly the reason that Lyman and Wes were talking about. That is that some of these take processing power.

The other issue I wanted to bring up, which is sort of the this would be useful for researchers. 002 – at least the words that are in it now – is not about stuff that is useful for researchers. And, in fact, what I believe most people believe now is that data from any individual RSO is probably useful for research. That is that you don't need to look across all of the RSOs. That RSOs are pretty typical.

So if we are considering things that, ooh, this would be really good for a researcher, individual RSOs might be the best source of that data. And it's not necessarily a requirement to do that across the entire RSS. Particularly because – again as Wes and Lyman said – this takes effort on the RSO's part to collect the 002 data.

KEN RENARD:

Yeah, good point. And that's where maybe we could use the optional phrase. If it goes into 002 at all. Lyman?

LYMAN CHAPIN: Paul, I would argue that RSSAC is not in a position to put requirements on the root server operators, period. But recommendations are fine.

PAUL HOFFMAN: That sounds perfectly reasonable to me as well. Please bring that up in the next discussion.

KEN RENARD: Brad?

BRAD VERD: Yeah, I just want to remind everybody about the first work party we did around 002. And there was a lot of discussion about research data. I mean we discussed that ad nauseum. To the point where we all agreed that we would not be including requests or data for research. This was about operations. So just to remind everybody.

KEN RENARD: Right thanks. Is there any objection to starting a work party to update 002 including not only considering label-count but looking at other metrics, new metrics or updating existing?  
Andrew.

ANDREW MCCONACHIE: Once again I'm not objecting. I'm just suggesting an action item for me, which is to send a statement of work around to the Caucus and ask for updates to the scope. So maybe we have a period of review for the statement of work where people can suggest new things for the scope. Discuss that in the Caucus. And then, once the statement of work is solidified, we kick off the work party. Does that make sense?

KEN RENARD: Sounds good.

ANDREW MCCONACHIE: Okay.

KEN RENARD: All right. Any other comments, suggestions on RSSAC002 before we go to 001? All right. With that, I'll hand it over to Paul for discussion of 001. Paul?

PAUL HOFFMAN: All right, thank you. I'll wait for the proposed statement of work to come up. Great, thank you.

---

So this came out of an idea that Duane Wessels and I had. And I will speak only for myself. Because Duane and I often disagree on things like this, although we did agree together to propose this.

So 002 gets updated approximately every two years. 001 has not been touched since it was originally published. And there are some things in there that we noticed are feeling quite stale. So I'll walk through the scope briefly first without giving any examples. But then I would assume, Ken, that the rest of the discussion would be is there anything to do here.

So the existing expectations in 001 may or may not be relevant to today's RSS environment. That is, they were written before a lot of what we were doing now—especially with the 038 work—came up. So go through each of the expectations that are listed in 001. And looking to see if they're relevant.

And then number two is an interesting one. Which is there are many things in 001 that say RSOs shall do this or will do this. And it's not clear that RSOs are doing that. Or if we even want that wording, particularly given what Brad just brought up. I'm sorry, what Lyman just brought up about requirements and such like that.

So we can evaluate whether the requirements wording is correct. Given that it may be that, in fact, many of these requirements that are currently in 001 aren't being met. And if they're not being met,



---

should we remove them? Which then goes into three, which is will/shall/may/must blah, blah, blah.

So two and three are linked in the sense that it is likely that some of the requirements are not being currently met. And therefore are they really requirements? Or should we be putting stronger requirements in? Most of these are around publication, that an RSO will publish a something. And maybe make it so that publication is easier and is easier to find.

So then, four and five are also similar. Some of the current service expectations might need to be revised — just again looking at the current world — and maybe new ones should be added or maybe not.

Which leads us into six, which is some of them might be removed. We now know more about how RSOs act, especially under requirements, than RSSAC did when it developed 001.

And then number seven, we know that some of the terminology is wrong. And, in fact, even though seven is the last one on this list, it was the one that actually bugged Duane and I during a particular phone call we had. Where we were like, “Oh, my God, it says that.” And so at a minimum we can clear up the terminology to use the current lexicon and such like that.

So with that, Ken, I’ll hand it back. But basically, this is why we were thinking that maybe opening this up would be good. But

---

there are worm cans here, particularly about the idea of requirements.

KEN RENARD:

Okay, thank you, Paul. Open it up any comments, thoughts, on what Paul said. In the room? Okay. Lyman's an old hand, okay.

So I feel like this is certainly worth a work party. Any document that's what now — is it five or seven years old — certainly worthwhile. I guess there's a chance that some of this could be overcome by governance structure. But honestly an updated version of this could be used as input to that governance structure. So I think that's certainly worthwhile. Wes?

WES HARDAKER:

Thank you. I think this is a no duh. But we do need to make sure that we correlate with the RFC on the other side. We've always been careful about the split of operational metrics are more on the RFC and these are a different set. So I think it's fine to have the work party go forward. And it's always possible that work party does a lot of analysis and concludes a document doesn't need to be changed. And we don't do anything.

That's actually an acceptable thing. Although terminology itself may be worth republishing. But I just want to make sure we draw the distinction. And we might want to put it into the work party

---

statement of where that line is so that we set a clear scope boundary.

KEN RENARD: Like those are roughly technical versus non-technical expectations?

WES HARDAKER: I could not come up with that wording in that few seconds I had to think and I'm not going to try. But, yes.

KEN RENARD: Okay, yeah. I think that RFC is about the same age but... Lyman?

LYMAN CHAPIN: Yeah. As one of the authors of the RFC, yes, I think this is a good idea. Wes made a good point. I suggest — as you said, Wes — that we put some wording into the document here. But also for this work party to have a look at the RFC to see whether that also needs an update in the IETF arena. And suggest updates if they see a need for that.

WES HARDAKER: Yeah, agreed. I mean those of us in the IETF also — which is almost everyone — can also write the RFC update. And push it through. But I don't think it can go into the charter per se, no.

---

LYMAN CHAPIN: No, maybe not. But we can encourage people verbally. Thanks.

KEN RENARD: Paul?

PAUL HOFFMAN: So that's actually I think very relevant and it should be in the Charter is that a review of 7720. Not that we can change it directly. Wes's point that anyone can write that, yes, but normally since 7720 — hang on a second, I'm pretty sure — was an IAB document.

Yeah, so the draft that led to 7720 was actually in the IAB string. So we could do something where if, during our review we say 7720 should be updated in this way, we could pass that to the IAB, either through the ICANN liaison or probably much more informally through, say, Wes. But not only is it not our capability to force changes into an RFC. It is definitely not our capability to force changes into an RFC that was published by the IAB. However, I would hope that the IAB would be interested in our input for 7720.

KEN RENARD: Thanks, Paul. My reading, my understanding is that these updates to each could be independent. Update to one does not

---

require updates to the other. But if it makes sense, we could nudge, nudge, hey take a look at that.

All right. So it sounds like we have pretty good consensus on starting up work parties for both of these. And some good input to the SOW. Andrew?

ANDREW MCCONACHIE: Is it the same for RSSAC001v2 where I should just send the statement of work to the Caucus and let people chew on it for a couple weeks? And then with adds and additions and then after that we kick off the work party?

KEN RENARD: Yep. That sounds good to me. Anybody?

ANDREW MCCONACHIE: Okay. Perfect.

KEN RENARD: Thumbs up from Lyman. Any other thoughts on this? Or anybody want to dive into some of the specifics? I don't want to put you on the spot, Paul. But if you wanted to talk about some of the specific things that you thought might be updated, you're welcome to but not required.

---

PAUL HOFFMAN: I would rather not. Simply because that's like I have desires and Duane has desires. And I'm pretty sure we actually don't agree. So I don't think me listing mine is at all relevant at this point.

KEN RENARD: Okay, thank you. Well, we've got both of those documents for work parties. And statements of works that will be sent to the Caucus mail list.

And if there's no other discussion on that, I think we could adjourn early. All right. Ozan, anything to close out?

OZAN SAHIN: Well, Ken. Thank you, everyone, for joining this session. The session is adjourned. Have a good rest of your day.

KEN RENARD: Thanks, everyone.

OZAN SAHIN: Please stop the recording.

**[END OF TRANSCRIPTION]**