

---

[DUANE WESSELS]: Welcome back to the afternoon session of RSS Metrics. Per the agenda, in this session, we wanted to talk about RSS correctness thresholds, I believe. We've already talked about the two correctness metrics and agreed to combine them into one. What we haven't talked about yet is thresholds, I believe, for either of the cases for correctness. Is that right? Did we – well, I forget.

[PAUL VIXIE]: Well, we talked about them. I don't think we settled—

[DUANE WESSELS]: We didn't settle. So this may be a very brief discussion because, in all the things that I've heard and in all the inputs to the spreadsheet, everyone is saying that correctness should be 100%, which is fine. I have no reason to believe it would be any different from the RSS versus the RSO metrics.

I will point out that the RSS correctness metrics has proposed are a simple aggregation of the individual RSO metrics. So I don't think we need to consider doing anything more complicated as we have done for the availability and latency metrics. This can be a straight aggregation.

[PAUL VIXIE]: Perhaps we should ask if anyone sees a need to have anything different. I don't, but let's open the floor if anybody has any thoughts on that.

---

*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.*

---

[DUANE WESSELS]:                   Okay. Doesn't seem like it. Let me ask the group. Can we take it as a given that we want to have 100% as the correctness threshold for RSS correctness?

Seeing some thumbs up and no dissent there. So that's essentially what I expected. I think that's good.

[FREF BAKER]:                   The only question I had was after our discussion this morning about how to combine the two metrics into one from the RSO perspective. Will that impact our current wording that we have in the RSS correctness matching? I'm not sure. I was taking a quick look and went, "[Hmm]." So it would be good if we did decide if we had to do some revision there. I'm not sure.

[DUANE WESSELS]:                   [inaudible]. I don't think so. I think it should just be the same. What it says in the document is that all the measurements covering a period of one day are aggregated together. The correctness is calculated as the number of secure responses – or, in this case, it would be correct responses – divided by the total number of responses received, expressed a percentage. So I think that's the right thing to do in this case.

UNIDENTIFIED MALE:               Good.

---

UNIDENTIFIED MALE:            Yeah.

[DUANE WESSELS]:            Assuming that we're done with the correctness thresholds for the RSS and the metrics, I think we can use the time in this slot to talk about some things that weren't originally on the agenda but that we want to talk about anyway. I have two things on my list. If anyone else can think of something that they want to talk about that's not here, please feel free to speak up.

The two that I have are – yesterday there was a suggestion about the RSS publication latency. Publication latency is the one metric that was omitted from the RSS section but remains in the RSO section. So that's a little bit of a discrepancy. That's fine. I think, again, going back in time, the discussion back then was that we had a hard time thinking about what publication latency for the system as a whole meant. But we can reevaluate that.

The other item I have is a suggestion to maybe omit the examples from this document. If you look at the end of this document, there's a longish appendix with a lot of examples. The suggestion was that those aren't maybe useful in this document. They might confuse the reader. We could publish them some other way in some other documents – not necessarily some RSSAC document, but publish them some other way if people would find them useful. So that is up for discussion as well.

If there's any topics that people want to bring up, now would be a good time to get them on the agenda. Otherwise, I think probably we'll be finishing a little bit early today.

---

Let's go back to the question about RSS publication latency. As I said, this is certainly something that we could add back. Some people might wonder why it appears in one section and not the other. If that's a question we're not prepared to answer and we want to include it, then we could certainly do that. We could come up with a method for calculating the publication latency of the entire root server system, again, perhaps as a straight average or a straight median of all the aggregated responses together. That would be an obvious first stab at this.

Does anyone have opinions about this question?

[PAUL VIXIE]:

Since I was the one that brought it up, I think it's important that we put it in so people don't ask questions about it. I don't think we have to spend much time on it. I think the straight average would be just fine. I suspect from our discussions today we're going to end up changing a fair amount of the RSS anything here, both is publicly report, what's internal for us, and such like that. But, as a first guess of just saying the averages, that, to me, is not a wrong to publish.

One thing that we are sure of, looking at the RSSAC 002 data, is that there is a wide dispersion of how quickly instances get updated even within one RSO. So none of this should feel like it's really the average. So just us reporting averages I think is just fine here.

[Touch it gently].

---

[DUANE WESSELS]: You said the RSSAC 002 data?

[PAUL VIXIE]: Yes.

[DUANE WESSELS]: Okay. Right. So just to note that that's 95% instead of median. So it may be a little more susceptible to outliers than this.

[PAUL VIXIE]: Yeah, although, again, looking at even just the stuff that I sent for Appendix B, one of the things that I saw was that, in fact, the current publication latency – oh, actually, no. That didn't get published here, but [the one] that I sent to the list. Even among instances seen by the eight vantage points, there was a fair amount of publication latency, even difference during the day. So one rollover happened in ten minutes for one, and it took 20 minutes. It's probably the same instance. So I think there's wide enough variety we can just say let's just publish as an average now and drill down deeper later.

UNIDENTIFIED MALE: Okay.

RUSS MUNDY: I thought what we were going to have for what got published was pass/fail against a certain threshold, rather than having numbers

---

published [as] an average. Is that different? Are we going to do this one differently?

UNIDENTIFIED MALE: This is for the ...

[DUANE WESSELS]: Yeah. That's another open question that we should probably discuss at this time. We should probably try to resolve that, not only for this but for all of the [RSS].

UNIDENTIFIED MALE: Yeah.

[DUANE WESSELS]: So let's get back to that in a minute. But, for now, let's stay focused on whether or not we do RSS publication latency.

UNIDENTIFIED MALE: [inaudible]

[DUANE WESSELS]: Yeah.

---

UNIDENTIFIED MALE: It seems like we should do it just for consistency with the root servers, or ...

[DUANE WESSELS]: I'm not hearing any arguments for keeping it out at this time. So, when we go back and work on the document, we'll put it back in. But that'll be the last time we put it back in. If it ...

[BRAD VERD]: If, when you work through the document, there really just isn't a need for it, just leave it in. They'll put a statement as to why we're not really addressing it. So there's a section in the paper no matter what, but you explain correctness. All the RSOs are correct, so therefore the RSS is a correct/right type of thing.

[PAUL VIXIE]: I think this will be closer to the RSS latency, which is—

[BRAD VERD]: [inaudible]

[PAUL VIXIE]: Yeah. But RSS latency is we're just going to do an average because that's what you all might want to say, not like that that's where we're going to set our own internal goals.

---

[DUANE WESSELS]: All right. Thank you. I guess let's move on to the next topic, if I can. The proposal is to remove the examples from this document. Ozan, maybe you can scroll down to that appendix.

UNIDENTIFIED MALE: A and B.

[DUANE WESSELS]: Appendix A. Appendix B. I think it's up higher. It's the one with all the colored bars. Yeah. What you're looking at here is Appendix A. This is something that we'd used in previous work party discussions. I'm not suggesting that we keep this. I think this should be probably taken out. I don't think this is particularly useful for the final document.

The question that I do want to talk about is the example. Can you scroll up, please, Ozan?

[PAUL VIXIE]: [inaudible] one of the examples – oh, right. Sorry.

[DUANE WESSELS]: We're talking about the same thing, right?

[PAUL VIXIE]: Section 8.



---

[DUANE WESSELS]: Yes. Thank you. Section 8. I was wrong. They're not in the appendix, yeah. They're in a separate section. Here you can see that there's an example for every proposed metric. They look like this.

Do people think it's valuable to have these in this document or valuable to have them somewhere else? Or not valuable at all? Opinions, please.

[BRAD VERD]: I think they're fine to include.

[DUANE WESSELS]: It was you suggesting to remove them. What was it? Because you thought it could be confusing for people?

UNIDENTIFIED MALE: Yes. I think, if we have examples in here before the data is actually being published, people would have an idea of how the data is going to be published. But we don't actually know how to be published, so putting an example in here locks us down. If we end up publishing something different, it'll be like, "Oh, are you going to also do this or not?" whereas, since we're going to describe fairly well in the text what needs to be reported – not that we've completely decided that, especially for RSS – I don't think having examples is going to help a reader understand where this is going relative to doing it. And we don't know, for example, if we're going to do it as a table this way or not.

---

So I think it could have the negative result that people are going to expect us to do it this way, and we may end up not wanting to do it this way. We may want to do it another way.

[DUANE WESSELS]:           Yeah.

UNIDENTIFIED MALE:       [inaudible]

[DUANE WESSELS]:           Oh. Anand?

ANAND BUDDHDEV:           I was just going to say these are just examples. It's fine. If people see something else later, you just say these were examples. This is not how we said we would present them. It's just one way of presenting results.

[BRAD VERD]:                I feel like a broken record. Just because it's in here, if we need to add verbiage that says, "This is how we're thinking about it now. It's susceptible to change. It's whatever because we don't have data," then put that in there.

UNIDENTIFIED MALE:        [Or just remove it].

[DUANE WESSELS]: Right. I think my preference would be to, as you say, add some text and be a little bit more specific about why these are here and the way we consider them examples so that, as someone gains implementation experience, they may want to present in a different way.

[BRAD VERD]: They were put here for a reason, yes? I assume. Somebody wanted to see it. Somebody was interested. Somebody thought, "This might be helpful."

[DUANE WESSELS]: Yeah, probably me. But what I found was very helpful was when we were looking at the IANA SLA reports. I like the way those were presented, so I thought to apply some of those—

[BRAD VERD]: I agree. I think many people learn and process data in different ways. Some people do it via words and reading. Other people do it visually.

[RUSS MUNDY]: So it sounds like add words and leave the examples in. Sounds like we reached ... okay.

---

[DUANE WESSELS]: So let's then get back to a question which we briefly talked about a minute ago, which was a proposal for, in the case of RSOs, that the reports would have a pass/fail indication. But for the RSS, the reports would have actual measured values. That's different than what's currently in the document. Currently it says pass/fail in both cases. So the proposal is to change it for the RSS and put in the reports the actual values.

Are people comfortable with this idea?

Is anyone not comfortable with this idea.

[PAUL VIXIE]: I guess we all need to make sure everybody has a good understanding of how the values get established.

[DUANE WESSELS]: Mm-hmm.

[PAUL VIXIE]: So that, when those values get published, each and every RSO-associated person can provide a description as to how we got there and nobody is embarrassed by what they see.

[BRAD VERD]: I don't want to rehash and of this, but I will just say, if they're different – I'll say this for the sake of Fred; I keep pointing to Fred because Fred is

---

going to be Chair next year – Fred needs to explain why they're different.

[DUANE WESSELS]: You mean why it's different for ourselves than RSS?

Yeah. I think we can add a reason to the document about how we came to this conclusion. I don't have a problem with that.

We're going to finish earlier than he and I predicted. One more—

DANIEL MIGAULT: Just ...

[DUANE WESSELS]: Go ahead, Daniel.

DANIEL MIGAULT: Just to clarify, here in this we're going to explain what is on the whiteboard? That's what we're talking about?

[DUANE WESSELS]: Yes. What's on the whiteboard is an RSS metric formula. So this new formula needs to go into the document. Then in the reporting we would include the actual output of this formula rather than a pass/fail for the threshold. We will still have thresholds and we can still say whether it passes or fails. But we will also publish the measured values or we will recommend that they published.

---

All right. One last thing that I just thought of from yesterday that we talked about in the context of availability was days versus months. Currently, in the document, all of the aggregated metrics are on days. So, every day, you calculate the metric and compare to the thresholds.

We talked about how, when we think about availability, monthly availability makes more sense for some people and that maybe we want to add a monthly availability threshold as well, or change it to monthly. I welcome opinions on this. Should all of the metrics be changed to monthly instead of daily? Do we want to add monthly? Do we want to not do it at this time? Lots of choices here. Opinions?

[BRAD VERD]:

From an operational perspective, I like the daily so that I can see what's going on and what the health of it is fairly regularly or as near real-time as possible from a reporting/talking to the world. It seems like monthly is more normal. So it'd be nice to have both. I don't know how hard that is.

[DUANE WESSELS]:

I think mechanically it's easy. I think the harder part would be, if we're going to have both, that we need to think about if the thresholds are the same for both or are they different. Because that's really where this came up in availability. We said we might want to have an availability threshold where, if you have a problem for a day, it doesn't affect your monthly availability, for example.

---

[BRAD VERD]: Why would they be different? I don't know. I'm trying to think through.

PAUL VIXIE: The easy case on why they would be different would be availability. That is that the monthly threshold is going to obviously have to be tougher than the daily threshold or else what you're saying was, in order to fail for the month, you have to have failed every day. So coming up with that greater toughness – that is, let's say that you fail one day on availability and that's it for the rest of the month. If we want a reasonable threshold for the month, we would want a threshold where you didn't have to fail every day in order to fail that threshold for the month. So we'll have to come up with a different, stronger number for the month if we're averaging the number of days in any way. We can do that.

By the way, I agree that having both would be good for the metrics, but I'm not sure how we would show a line from the threshold for daily that we've chosen to the threshold for the month.

[BRAD VERD]: You have the equation for daily, right?

PAUL VIXIE: Yeah. So how we would aggregate those?

[BRAD VERD]: [Times the] number of days of the month. Do the same math.

PAUL VIXIE:                      Okay. So when you're saying "monthly," you don't mean an aggregation of the dailies. You mean an aggregation of the data over the month.

[BRAD VERD]:                      Right.

PAUL VIXIE:                      Okay. Sorry. I didn't get that. Yes. Great.

[BRAD VERD]:                      [inaudible]

PAUL VIXIE:                      Oh, absolutely. Absolutely.

[RUSS MUNDY]:                      Would this just apply for availability, or do we want to examine the other metrics?

UNIDENTIFIED MALE:                      [inaudible]



---

[RUSS MUNDY]: I mean, the general question was across all the metrics, but it sounds like we focused back to availability again. So it sounds like there's a strong desire for availability data and metrics for day and month.

[BRAD VERD]: Let me ask a different question then. If we apply what I thought was monthly, like all the data applied for the month, to all the metrics, does that add any complication or need to come up with a different threshold?

[PAUL VIXIE]: My view is availability and latency are similar enough that way to a common understanding. Correctness 100% is 100% regardless. Publication latency, in my mind, to most people will feel like response latency, which is, "Smaller numbers are good. I don't really understand your numbers. They don't really affect me." So doing it across the month is just as good as well.

The funny thing, though, as I was saying to Fred earlier, is that doing publication latency over the month is going to give us different values than doing it by day because, by day, you're starting in the middle of a publication stream.

[BRAD VERD]: Yeah. [inaudible]

---

[PAUL VIXIE]: Well, no. Actually over a month is a better way to measure it than over a day.

[BRAD VERD]: [inaudible]

[PAUL VIXIE]: Yes. Publication latency over a month is a more valid way of measuring than over a day.

[BRAD VERD]: Because you [inaudible].

[PAUL VIXIE]: Right. And your breaks of when did and SOA change make much more sense. I would almost argue then, if we're doing publication latency over a month, we don't want to do it over a day. But people are then going to say, "Wait. Why aren't you doing it over a day?" And just the numbers will be different and la, la, la. I think that'll be fine.

[BRAD VERD]: [inaudible]

[PAUL VIXIE]: Yeah.

---

[BRAD VERD]: And make the call [inaudible]

[PAUL VIXIE]: Right.

[BRAD VERD]: [inaudible]

[PAUL VIXIE]: No.

[BRAD VERD]: So it's a lot of work.

[PAUL VIXIE]: So, having done it now, it will take literally exactly the same amount of work to do month versus day and such like that. The central reporter that I'm using, which is just a single-core VM somewhere, crunches this stuff, like the daily numbers, in less than five minutes. So really this is not going to cause a sweat for anybody.

[DUANE WESSELS]: I'm thinking back to the example that Liman was using yesterday about [fingering] the routes or something like that, where you're down for 15 minutes. My concern is that, if you have two measurement intervals – month and day – and if you don't choose the thresholds carefully, you

---

can have weird situations where you can fail in 14 days out of a month but still pass the whole month, right?

LARS-JOHAN LIMAN: [inaudible]. I agree that is to be avoided, yes.

[DUANE WESSELS]: So I think there are situations where you may want to have a different threshold for a day than you have for a month.

LARS-JOHAN LIMAN: Yeah, I agree to that. Yes.

PAUL VIXIE: I'm not sure if I agree to that because any threshold that you fail – Brad, you're going to love this; just warning you – would make you get thrown out of the root server system, even if it's on the daily ... So, I know, going to that [extreme]. If we have a chart for days and a charts of months and there's a whole bunch of reds on the chart for day and none on the month, somebody looking at that is still going to totally understand that you didn't do it right. We're talking about adding month as a second measurement, not as an exchange for the day measurement. So with your extreme of "You have failed less than half the time but you still pass," someone is going to still see that you failed half the time. I'm not worried about changing that.

---

[DUANE WESSELS]: They certainly will see it, yes. I think the concern that I had is someone may say, “How’s it possible that you can fail have of the days in the month but still pass on a monthly basis?”

UNIDENTIFIED MALE: Math.

UNIDENTIFIED MALE: Right.

LARS-JOHAN LIMAN: Sorry. that argument will work on exactly 0.0001% of the population.

[NAELA SARRAS]: We are. If you were to apply this logic, we wouldn’t display the red. We don’t do the daily. We do a monthly threshold, right? So there are things that do fail, but it’s within the threshold that they don’t appear in the report, right? So, yes.

I imagine the CSC, if they wanted to dig more into the data and say, “Give me what failed the threshold,” then we would still present that data. But you’re right: they don’t see it. In the real-time dashboard, it doesn’t appear.

UNIDENTIFIED FEMALE: Yeah, I think you should go with monthly.

---

LARS-JOHAN LIMAN: I value that experience, so maybe we should listen to Naela.

[FRED BAKER]: Again, I think we're having a hard time keeping straight what we're going to report publicly and what we're going to be able to know ourselves for RSSAC and RSOs. If the RSOs need outputs daily, then that's something that I think can be communicated in some manner, perhaps independent of the metrics document, and we report monthly if we think a month's cycle is better for reporting.

[DUANE WESSELS]: Yes. We could maybe also argue that the operators needs something other than monthly or what's in the document. They can go and do it themselves, that it's not this work party's job to solve this problem. They have the ability to do their own monitoring if they so desire.

LARS-JOHAN LIMAN: I still argue that the underlying data should be public, which means that anyone can do this – crunch the numbers – including the root server operators, and produce whatever reports they see fit. But we don't necessarily need to hold it up in the data report.

[FRED BAKER]: So are we closing in on an agreement of monthly-reported outputs for the public publication? Is that what we seem to be leaning towards? And daily? If people want to, they can get the data? Because I think we agreed earlier that the data would be public.

---

Okay.

LARS-JOHAN LIMAN: I support that, and I would add as a comment that, if it turns out in the future that people actually ask for a denser schedule, then we can change.

[FRED BAKER]: Very good point.

[DUANE WESSELS]: Yes, that is different. But I think monthly has certain nice properties relating to people's attention spans and things like that. It's easier to think of it in terms of monthly reporting, whereas a lot of people aren't going to go to the website every day and check what the latest is.

NAELA SARRAS: Duane?

DUANE WESSELS: Yes, Naela?

NAELA SARRAS: Also, just speaking from experience, I think, for however oversight is established over this, I think monthly makes a much better cadence for your oversight to look over it and ask questions and request more data.

---

You're not going to have an oversight looking at it every day. That's just a waste of everybody's time. Thank you.

[PAUL VIXIE]: Naela, I have a question for you.

NAELA SARRAS: Yes?

[PAUL VIXIE]: By monthly, we actually mean from midnight on the first of each month, which has a different number of days. It's not like 30-day periods with five days stuck at the end.

Okay, great.

NAELA SARRAS: Yeah. That's exactly how the dashboard works and UTC, just like you guys talked about yesterday.

[PAUL VIXIE]: Okay.

[DUANE WESSELS]: So what I'm finding is, the longer that we talk, the more work it is for me and others who have to [write it] So I'm motivated to wrap this up pretty quickly, or, of course, solicit help in writing the document.



---

[DANIEL MIGAULT]: [inaudible]. I'm just looking at the whiteboard. I think, while you have curves with eight ... So I think eight is a threshold, so the curve is actually designed with this threshold. So that's only one comment: that we should ...

[DUANE WESSELS]: Yeah. Certainly, we'll take this concept and put it into the paper appropriately. It won't go in verbatim. We'll run it by everyone to make sure that everyone still has the same understanding of what we discussed here in our meeting.

I really don't have any other topics.

UNIDENTIFIED MALE: [inaudible]

[DUANE WESSELS]: I did.

UNIDENTIFIED MALE: There's a mail.

[DUANE WESSELS]: There's a mail? All right. Are you talking about – which one are you talking about?

---

UNIDENTIFIED MALE: [inaudible]

[FRED BAKER]: That he just sent this morning.

[DUANE WESSELS]: Right. That was not sent to the whole group, but we can talk about it.  
Can I connect my laptop to the screen?

[FRED BAKER]: I don't know why I was ...

[DUANE WESSELS]: Well, how many remote participants do we have besides Shinta?

UNIDENTIFIED MALE: Three.

[DUANE WESSELS]: Three? I guess you can do it. I'll tell you what to type.

UNIDENTIFIED MALE: [inaudible]

---

[DUANE WESSELS]: The background here is that myself and a couple of others received an e-mail from Shinta mentioning some concerns about response latency thresholds. As a reminder, we agreed that, for [EDP], for example, the latency threshold would be 250 milliseconds. Shinta was looking at data from the RIPE/DNSmon to see how well that matches reality, I guess.

Can you go to that website, Ozan? It's [atlas.ripe.net/dnsmon](https://atlas.ripe.net/dnsmon).

[PAUL VIXIE]: M-O-N.

[DUANE WESSELS]: Okay, thanks. Under the grey box, where it says, "DNS responses for root," there's another one that says, "Unanswered queries." Can you click on where it says "Unanswered queries"?

UNIDENTIFIED MALE: Right below that.

[DUANE WESSELS]: Down a little bit lower. Select "Response time." Okay. Can you – I guess we're going to zoom in on ... Let's zoom in on B Root, the third line down. That's not what I meant. Just click on it.

[PAUL VIXIE]: Just click where you are. Yeah. Okay.

---

[DUANE WESSELS]: Now if could click on the—

UNIDENTIFIED MALE: It's actually still thinking.

[DUANE WESSELS]: Well, it might think for a long time. I think you can still do it.

UNIDENTIFIED MALES: [inaudible]

[DUANE WESSELS]: See at the top of the page there? There's a little green box and a little red box. Can you click on that area?

UNIDENTIFIED MALE: [inaudible]

[DUANE WESSELS]: Yeah.

UNIDENTIFIED MALE: [inaudible]

---

[DUANE WESSELS]: Yeah. All right. Now what I would do here, Ozan, is – where it says 60 in the green box, can you change that to 250?

UNIDENTIFIED MALE: 250.

[DUANE WESSELS]: 2-5-0. Okay. My interpretation of this is now that all this data points, if they are less than 250 milliseconds, are in green. If they're over, they're in red. I some red, but I see more green than red, which to me means that this server would meet the median threshold of 250 milliseconds response time.

[BRAD VERD] Yeah. [inaudible]

[DUANE WESSELS]: Mic.

[BRAD VERD]: If you can go back to the previous page where it just shows the roots ... And now there's an even better [inaudible].

[DUANE WESSELS]: Yeah. Again, by doing the things that [inaudible] to do, clicking on here (the page) would show any root server at any time. I guess this may be averaged. I don't know. If, average response time is over 250, it would

---

appear in red. So at all times, for this day, anyway, all the root servers are less than 250.

[PAUL VIXIE]: Are these medians? Or are these 95% percentiles? Do we know?

BRAD VERD: Don't know.

[DUANE WESSELS]: I don't know if it's—

BRAD VERD: [Not trying again].

UNIDENTIFIED MALE: I think they were averaged.

[DUANE WESSELS]: They might be mean instead of median. I don't know.

UNIDENTIFIED MALE: [inaudible] green boxes.

UNIDENTIFIED MALE: Just hover over any green box.

---

UNIDENTIFIED MALE: Median.

UNIDENTIFIED MALE: Right.

[DUANE WESSELS]: So this is very close to our metric as we're proposing it. This is medians of response times. I think we're well above 250. You could slide the data back in forth in time or zoom in or out if you want to do, but I expect that it's going to be almost all green.

BRAD VERD: I just wanted to address it since it was brought up.

[DUANE WESSELS]: Thank you, Shinta, for the e-mail, but I'm convinced that 250 milliseconds is a pretty generous threshold. I've looked at other RIPE ALTAS data and came to the same conclusion.

UNIDENTIFIED MALE [inaudible]

[DUANE WESSELS]: Yeah.

UNIDENTIFIED MALE: [I'd say] overly generous, I think.

UNIDENTIFIED MALE: [inaudible]

[DUANE WESSELS]: One last thing that I want to talk about was, during the lunch break, we had a little meeting, Russ and I – Fred was there, and Steve – talking about our thoughts on the upcoming schedule for how to get this work to competition and the next milestones and so on. We've been having work party calls every two weeks, so there'll be another one in about two weeks for now that's scheduled for October 17<sup>th</sup>. Obviously prior to that work party call there's a lot of work for us to do to get the document into shape and add new things and change things and so on, so we'll be doing that. By the time of that work party call, we'll try to have the document in pretty good shape. The goal of that call will be to present to the whole work party or to people who weren't here the things that we discussed today and the new recommendations and changes.

Two weeks after that, there would be another work party metrics call. Normally, that would be scheduled for October 31<sup>st</sup>. However, that's also when the DNS-OARC workshop is occurring. So we're going to propose that a couple days earlier to Tuesday, October 29<sup>th</sup>, for the work party call. Our goal at that point would be to have a first final draft of this document.



---

Then, at the ICANN meeting in Montreal in early November, we hope to take some time on the agenda of the caucus meeting to present an update on this document, probably very similar to what we opened this meeting with. That slide deck will probably be the basis for that. Maybe spend an hour presenting the work to the caucus members.

Then, not too long after that will be the caucus meeting at the IETF in Singapore. We'll do something very similar. We'll present the status of this document to that group.

In late November is our tentative deadline for having the RSSAC Caucus provide its final feedback and input to the document. In December, we would deliver a final document to RSSAC for its final comments and eventual voting.

If we're able to meeting this schedule, again, RSSAC would see this in December and maybe even have a vote in January. I wouldn't want to place bets on that actually happening, but that's the plan at this point.

Any last minute comments before we wrap it up for the day? Or is there any other RSSAC business that we need to talk about? We do have some time on the agenda tomorrow to wrap up the metrics work, but I want to save that for tomorrow.

UNIDENTIFIED MALE: Recommendations.

---

[DUANE WESSELS]:                    Yeah, recommendations. If there's nothing else today, I would like to be done with the metrics work for today.

UNIDENTIFIED MALE:                "I would like to be done."

[DUANE WESSELS]:                I would like to be done.

BRAD VERD:                         Do we have anything else [inaudible]?

UNIDENTIFIED MALE:                [inaudible]

BRAD VERD:                         Fred? Well, hold on. Anything, Fred, that we should cover?

FRED BAKER:                        The one thing that comes quickly to mind – I mentioned this in the meeting [inaudible] – is that I know everybody here was really looking forward to a two-week workshop in May. But I'm thinking that, if we can actually line this up, we'll probably be [going to move] the spring workshop.

---

UNIDENTIFIED MALE: We are going to or not going to?

FRED BAKER: I think we avoid the need. To me, that's a good thing.

BRAD VERD: I don't know if it's a good thing. I am in favor of no workshop, but this group gets worked on face-to-face.

[FRED BAKER]: Okay.

BRAD VERD: So I'll just leave that out there.

FRED BAKER: We may think of a reason then to have a workshop. We'll make sure it's at least two weeks.

LARS-JOHAN LIMAN: At that point in time, we might already have a cross-community working group going on RSSAC 037 and the evolution [of that].

BRAD VERD: But GWG will be in effect by then.

---

LARS-JOHAN LIMAN: Yes. That could generate things for us to discuss. So please don't rule it out entirely. But I realize that we are planning the horizon here, especially for staff to help us with that. But, yeah, don't just drop it, please.

We'll give time back.

UNIDENTIFIED MALE: [inaudible] going just to [inaudible]

LARS-JOHAN LIMAN: Yes.

**[END OF TRANSCRIPTION]**