

---

FRED BAKER: Okay, so Duane and Russ, do you want to discuss the recommendations in the metrics document?

DUANE WESSELS: Yes we do, well, we said we would. This is the last session for metrics and we have two things that we wanted to go over; one is the section of the document that talks about recommendations and then I thought we would do like just a summary wrap-up of the last couple days and what we've all talked about. So, actually, Ozan, can you put up the Google doc first? Go to Section 7, I think. Alright, I'm going to have to look at my laptop here.

FRED BAKER: And the reason for our relocation to table, Duane noted yesterday, quite appropriately, that there were times that we on the side that we missed some people further down, so we want to be able to see when people raise their hand and ask to give a contribution. Sorry about not being better positioned earlier.

DUANE WESSELS: So, if you're following along, we're going to go through Section 7 of the draft document and there are a few recommendations here. Please feel free to jump in with comments or questions at any time. So, the first recommendation is sort of about, what in the document we call "the official implementation of the metrics system," it says monitoring

---

*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.*

---

system here, but it probably should say metrics, because this is not about monitoring, and there are a few sub-bullets.

So, for example, the first bullet says that it should "meet the minimum requirements specified in Section 3 regarding the number location and connectivity and other requirements of the vantage points. So again, this is designed to apply to the official implementation of this. You can interpret that to mean whatever the PMMF or its derivative comes to rely on additionally, there is a recommendation that the implementation must be published as open source software so that people can review it and had trust in it and that it's transparent in that way.

RUSS MUNDY:

One question that I had from the current wording, it's not clear to me if the intention for having open source publication, if it's for primary just review and confidence in the software or if another aspect so that other activities can take the software and actually do somewhat of their own running of similar system. Do we want to be any more definitive as to reasoning associated with this recommendation?

DUANE WESSELS:

My inclination would be not to specify that. If I'm reading this as someone who has been tasked or funded to implement this, I'm not sure I would welcome seeing that, because I might think that then I have to support those other uses, as well. I think this is sufficient, and certainly an open source license would allow someone to take that software and use it under the terms of the license. I think the main

---

intent is to make it transparent. That's my personal opinion, I don't know if anyone else feels differently.

So, the third bullet says, "Must openly publish daily data on when individual" - it's changing in front of my eyes, "Must openly publish data on when individual RSOs are passing or failing any threshold. So that's pretty straightforward.

PAUL HOFFMAN: Duane?

DUANE WESSELS: Yes, Paul.

PAUL HOFFMAN: But I think we're going to have to change that based on the discussion yesterday about RSS and RSSAC getting data. "Publish" was a clear word a month ago here, and now it's not, so I think that will have to get changed to make it clear that the data also has to be available to RSSAC in some way.

DUANE WESSELS: Yes, I agree. The discussions that we've had the last couple days are going to change some of these recommendations, and this one in particular, because it talks about publishing data which is something we've talked about. So, while this bullet talks specifically about the RSOs, yeah, there will be changes here. And the next bullet in my mind

---

is sort of similar, because it talks about pass/fail indication for each RSO. Here it says, "should only publish pass/fail indicators for each RSO and not the actual measurements or metrics."

And again, I think this is something that might change because we may want to be more clear about the reporting that comes out of this versus the raw data that may also be available to either certain parties or to the public. But I think you can kind of get the intention here, again, these are targeted towards the official implementation or the official reporting for this and we do note here in this bullet, one of the reasons is to discourage gaming the system.

RUSS MUNDY:

I'd like to get confirmation from the work party that in fact as this work party we feel we want to have all of the data that is collected as part of the official system be available and do we want to include that as part of the recommendation? What are people's thoughts on that? First, do we still feel all the data that's collected should be publicly available? Yes? No? Because that was the last two days discussion led to that.

BRAD VERD:

I don't think it was a statement saying it should be publicly available, the statement was if it's not publicly available you need to have a reason why, and you've got to have some justification. I think as a result of that statement, I think Fred said it should just be publicly available. If I'm Fred, I don't want to answer that question.

---

DUANE WESSELS: That kind of matches my sentiment of the discussion yesterday as well. So again, this will probably change, this recommendation will change to reflect that sentiment.

FRED BAKER: If you're looking for me, too, I also agree it should be publicly available.

RUSS MUNDY: This was kind of the sense I was wanting to get and the number two question is related to that, do we want to put that in this document? We're not required to, but we can. Are there drawbacks either way? If we should write it into here or just leave it unspoken, but our work party agreement is that it will be. Thoughts on that. Paul?

PAUL HOFFMAN: I think it absolutely should be in the document for two reasons. One is that since this document is partially meant as input to 03-7038, when the PMMF or whatever it's called comes around, they should know that the data is also going to be used for that, that should be something that they know, but more importantly, if someone reads this document and doesn't see that and then is surprised that, oh wait, I wanted that data, I didn't know it was available, no one told me it was available, that would be a negative surprise, as well.

RUSS MUNDY: Okay, Paul's suggestion, should we get it in the recommendation? Anybody opposed or different thoughts on it? Fred?

---

FRED BAKER: I would suggest that it be a separate recommendation, that it would be highlighted as this should be true.

RUSS MUNDY: A separate recommendation, okay.

DUANE WESSELS: Okay, so the next bullet, I guess it's #5, says that "the official implementation should publicly describe its methods for collecting measurements and aggregating metrics including the topological location of each measurement vantage point. This description should be complete enough for RSOs and DNS researches to create their own measurement collection systems similar to that used by the PMMF.

So I think you see the intention here is that if people want to build their own systems that would get similar results out, that they would have enough information to do so. The description of the official one would allow them to do that. I'm not sure we want to mention PMMF here but that's sort of the intention. Any comments about this one? And obviously if you're at all familiar with [inaudible] you will notice the capitalized Must and Should, we sort of adopted that terminology. Oh, I'm getting a dirty look from Paul.

PAUL HOFFMAN: I was going to wait until the end on this one, but that's terrible. I mean, just as one of the people who has been involved in the 2119 discussions

---

for about 25 years, just make everything a must, absolutely, but if this is a recommendation, every "should" that you expect anyone to do anything about needs to have, except in these cases, and we don't have that here, just say must, and you don't even have to capitalize it.

DUANE WESSELS: That was going to be my suggestion. Let's not necessarily adopt this terminology, especially without referencing it and what not, so make our lives simpler and lower case these words.

PAUL HOFFMAN: And make them all 'must.'

DUANE WESSELS: Okay, that's fine with me.

UNKNOWN SPEAKER: Do we agree on using non-capital letters and saying must instead should in capitals?

DUANE WESSELS: That's the proposal, yes. Is that okay with you?

UNKNOWN SPEAKER: Yes.

---

FRED BAKER: In that, then, do we need a reference to RC8119 or whatever?

PAUL HOFFMAN: No, we're not using those.

DUANE WESSELS: We're going to take those out. The next bullet, again, maybe sort of overcome by events, assuming that the raw data will be made public, this is sort of along the same lines, but here it says that the underlying measurements and metrics should be provided to RSOs, especially those that lead to failures of thresholds, failure to meet a threshold. It says the shared measurements and metrics must include all measurements from around the days of the failure and must include all measured values for all transports.

So again, if we're making the raw data available, that is probably sufficient. If there is something that is not going to be made available publicly then I agree with this recommendation. As much information as possible should be provided to them so that they can understand and investigate the times when they fail to meet thresholds.

RUSS MUNDY: One question that this raises in my mind, we spent some time in the last couple days talking about paragraph 4.8, which is the exceptional events paragraph, and that includes, I think it's a should, or a recommendation that additional data should be collected and made available for subsequent analysis. Do we want to say anything in the recommendation section about that data? Because I think we all



---

generally acknowledge that data is data that is sort of above and beyond, or somewhat different, or in addition to the normal collection data. So the weird events where we wanted to get the additional things, do we want to say anything about that data in the recommendation?

PAUL HOFFMAN:

I don't think we can until we see what the new text is.

DUANE WESSELS:

And then the last bullet here I think is probably just going to go away, because it says that the raw data would not be made public, but we here agree that is not going to be the case anymore, so that will go away. Alright, so that's Recommendation 1. Again, that's targeted at the official implementation.

Recommendation 2 says that, "RSSAC should begin investigating a better long-term plan for the location of vantage points." So this ties back to some in Section 3, I think. "Such a plan would distribute vantage points by networked apology instead of by geographic region." So, that's a future work item. I guess at some point RSSAC would take that and develop it into a more filled out statement of work and so on.

RUSS MUNDY:

One thing, I didn't know if we wanted to include in this recommendation or a different recommendation whether or not review of analysis approach that's described in various parts of the document should be itself undertaken in some amount of time, to see if both for

---

the values that are used for the thresholds and the techniques that are being used are the most appropriate that we can come up with, or if other improvements can be added at a later date.

DUANE WESSELS: Yeah, I think that would be perfectly appropriate to make a recommendation annually, or however often the whole thing should be reviewed. We do that on our other documents.

BRAD VERD: Right, we do that on our other documents, and then no one pays attention to them, and then we find out two years later that we said we should do it annually, and we're like, ooh, so quite honestly, somewhere in here they should just state that this is our first stab at this and we expect it to change and evolve over time.

DUANE WESSELS: Okay, that's a good way to put it. I'm not sure, we'll see if we can phrase that in terms of a recommendation, but I like those words. Okay, the next section is specific to RSO thresholds. My thinking is that we will also mention the threshold values in the actual metrics definition sort of at the end and probably also include the rationales there.

This part here, Section 7.1, may be just a summary of those, it may take the form of a table or something like that, but it would summarize the threshold values for each of the metrics and there is a section for the RSO thresholds and then a section for the RSS thresholds. There will be

---

a lot of revisions here because of course we've changed things, like we've combined the correctness metrics and so on, and we have actual values now based on the work party output. So this is going to change. I'm not exactly sure what final form it will take, but I plan on keeping these sections here and documenting the thresholds.

PAUL HOFFMAN:

This is Paul. So we had talked about documenting the rationale, although the rationale that I proposed for the RSO metrics, some of that changed, but some of it didn't, like the overarching, which one is more important. Do we want to have the overarching rationale here? If you're going to put thresholds in the main body, should we at least have the overarching rationale here? Do we want to put it somewhere else? This is sort of a where does it go in the document? But if you're putting thresholds early, then the rationale should be even earlier.

DUANE WESSELS:

I don't know. Speaking personally for myself, I'm comfortable with having rationale description for each metric and I guess I don't feel the need to have an overall overarching rationale, but that's just me. I want to hear from other people, if other people would like to have that, and if we can agree on text, then absolutely, let's include it. But again, you know my thoughts on that.

---

DANIEL MIGAULT: I don't understand the two ways. One is that in each section we have the rationale for each metric and the other way I understood to put that before or...?

DUANE WESSELS: We talked yesterday in one of the sessions about rationale a little bit, and Paul Hoffman had sent an email to the mailing list sort of outlining his rationale for how he came to his recommendations, and I asked him to add that to the end of this document. So if you scroll down to the bottom of the Google doc, you will see it there. I think what Paul is suggesting is take his writing and maybe move that early in the document so that a reader has a sense of the work party's overall rationale for how they came to think about these metrics and the thresholds. That's what I hear Paul suggesting. What I'm personally suggesting is that that's not necessary, but I'm not opposed to it, either.

DANIEL MIGAULT: I need to go through the document to see, but I will.

DUANE WESSELS: I think there will be more recommendations coming to the document, again, based on the work of the last couple days. We'll also add maybe a section here about future work. I'm not sure if it fits in recommendations or not, or if it stands on its own, but I do envision more recommendations coming out of this and some of these we can talk about when we, we're going to go through and summarize all the work from the last couple days, so some of that will show up there, too.

---

But if anyone at this point knows of other recommendations that should appear here that are not listed, please speak up and let us know. Naela?

NAELA SARRAS:

Thank you, Duane. I don't know, this is to think about, and maybe I'm bringing this in too late. So, this morning from the RSSAC meeting we just had and Liman's update on the CSC and the process that took us almost three years to put together a process to change the thresholds, I think there should be a recommendation, especially because of what Brad was saying earlier, this is our first stab at this, we don't know if these are the right thresholds, we need to see some operational data. And if this becomes part of contracts and you have to live up to these thresholds, I think you should think about a recommendation to think about a process to come together and change the thresholds should the need arise. Again, just speaking from personal experience on the CSC and the naming side.

DUANE WESSELS:

Okay, Liman, go ahead.

LARS-JOHAN LIMAN:

I can only wholeheartedly support that.

DUANE WESSELS:

It would be great if we could get you to contribute text, because absent that, I don't really know what to put, other than to say there should be a

---

work party. So if you're suggesting something other than a work party for a more efficient process, that would be good to have written down.

LARS-JOHAN LIMAN: I think the most important part here is that it makes sure that we have a process that is outside any contracts that are written, so that we don't have go through 15 contract parties and have every change to this set of metrics approved. So there is a separate process for modifying and changing it, which is more lightweight, so that you can sign the contract and say yes, we rely on this process if changes have to happen.

NAELA SARRAS: Agreed, and we can maybe provide maybe a little bit of text of what we think should go in there. Yeah, we'll work together.

LARS-JOHAN LIMAN: But you will have to remind me.

NAELA SARRAS: Yes, I will, of course.

DUANE WESSELS: That's good. I have a followup question. So this document at all does not consider contracts really in any way. Should it? Should it say like the work party recommends that if there are any contracts that reference the metrics, that those contracts should acknowledge that thresholds can change over time based on blah, blah, blah, and not

---

require contracted party, I don't know, am I getting into the danger zone here?

NAELA SARRAS:

Do you mind if I go there first? I'm not a lawyer, thank god. I don't think that's what it should be, I think the contract, if there is ever a contract, they should say you have to live up to these metrics that exist in this document. This is what we ended up doing with the naming contract. So, the metrics live on their own and the contracts refer to these metrics that we have to comply with. So the metrics can change under the oversight of the CSC and all the contract does is just defers to the most recent that we're complying with.

So I think it should say something like this, this recommendation, however this comes through, it should say there needs to be a process for changing the metrics, and then we develop that process. And then the contracts if they should ever come about, then they should just refer to that process and the metrics. I'm kind of making things up.

LARS-JOHAN LIMAN:

I have a different opinion here. I think saying contracts just goes beyond what this document is about. This is talking about metrics, and we're not presuming one outcome or the other for RSSAC 37. So where I acknowledge it's useful to have that conversation, but probably not here, but maybe of those entities forming those entities, there is a work item there, I think that's the most natural place. So to me, it seems we're jumping ahead too much. Thanks.

PAUL HOFFMAN:

A way to get what I think Naela and Liman want simply is to have, it wouldn't be a recommendation, but a very firm statement in the document saying that the metrics in this document may be updated by additional documents in the RSSAC document series. You don't even have to say if this one is going to be RSSAC 050, that it's RSSAC 050 V1, you can just say may be updated by additional documents. So it ignores the contracts. The contracts then go into what Naela just said, can say must follow, any contract must follow RSSAC whatever, or its successors for metrics.

Naela, I don't know if it was clear to everyone else, but it took me a bit to understand that you were actually not only talking about contracts like for root server operators, but for like TLDs and such, where we wrote in metrics in the TLD agreements, and then wanted to update them and they said we can't update them, this is going to be too hard to get new contracts, at least that's what I've heard within ICANN.

NAELA SARRAS:

I think Fred is trying to get in. No, I wasn't even referring to those. All I was saying is IANA has to perform to thresholds that were defined by the community during the transition and then we discovered that the thresholds were not sent correctly and then we had to invent the process to change those thresholds.



---

PAUL HOFFMAN: So, anyone in here who is also a TLD operator is familiar with the fact that the first set of contracts in the new GTLD program had some required metrics that didn't pass the sanity sniff test initial, but because rewriting contracts was so difficult, it would have been better if those had said point over here and let that happen.

RUSS MUNDY: One of the things that I wanted to add to this discussion is the part that we were talking about earlier about revisions and changes for the future, would likely be a good place to target getting these words into. So I'll edit the note I took earlier and see that we add that in.

DUANE WESSELS: I believe that's it for Recommendations. So, should we move onto the next part? Summarizing the days? Alright, Ozan, can you put those slides up? So, I've got a little slide deck here, the purpose again is just to reviewed what we've talked about and have something sort of in writing so that we all agree to what we agreed too. So we'll go through these one by one.

One of the things we talked about earlier was root server system thresholds and we agreed that the work party should recommend thresholds for RSS metrics. We need to important the purpose section to clarify that this document has purposes other than just RSSAC 037, and as an example if there is not, or until there is such a governance system, RSSAC may want to be notified of certain RSS problems and maybe take action on those. Okay, next.

---

We talked about the need for two types of thresholds, minimum and good thresholds, and we agreed that at this time the work party would only make recommendation on minimum thresholds and we'll add some text to the document that other probably higher thresholds may apply on a case by case basis for root server operators that enter in agreements or where there is an exchange of money for service. So, jump in quickly if you have comments, because I'm just going to kind of go through these, otherwise. Next.

We spent a lot of time talking about how we came to the idea that there is a certain number of root server operators that need to be up in order to provide acceptable RSS service and this is kind of the formula we agreed on. So  $\frac{2}{3} N$  minus 1, rounded up. That's what that formula means. So for 13 servers we end up with 8 as our value of K. So the rationale for this is that when one of the root servers is down a second query will be successful with two thirds probability. That matches everyone's recollection? Yes? Okay. Oh and by the way, my thought on this particular item is that we'll add this text and this formula probably to the top of the RSS metrics section because a number of the metrics sort of now depend on this value of K, so that's probably where that would fit in.

PAUL HOFFMAN:

This is Paul, except we also use this idea in coming up with the RSO availability metric and threshold. So it might need to go earlier or if not, there will have to be a forward reference, and I just generally don't like forward references. So consider maybe putting it earlier, as well.

DUANE WESSELS: Okay, yeah, I was thinking about where it would fit best, and I thought it would fit best there, so yeah, we'll think about that, if it makes sense or not, you're right. Okay, next one, please.

So we talked a little bit about this briefly. I think currently the document is a little bit underspecified in this regard on how measurements from bad vantage points should be discarded. So I know this is something I've talked about with my colleague, Matt Weinberg and he's been volun-told to write some text around this, so we'll have a proposal for this and firm this up a little bit. Next.

At the end of yesterday we had this discussion about daily versus monthly and we all agreed that monthly makes a lot more sense, both in terms of the way people would consume the metrics and how we like to think of it. So we'll be changing that throughout so that all of the aggregated metrics are reported on a monthly basis. Robert?

ROBERT STORY: I agree that publication is good for monthly, it will be available to the RSOs for daily aggregation so we can maybe act on issues more quickly.

DUANE WESSELS: I don't know, I have two conflicting opinions about that. One is that as we said yesterday, an RSO would certainly be free to do their own types of measurements if they felt the need to act more quickly. On the other hand, it's probably not a lot of extra work for someone to do it both ways, so I don't know, what do other people think?

---

ROBERT STORY: Well, I'm sure all the RSOs do their own monitoring but other vantage points see other things that we don't see, it would be useful to know.

RUSS MUNDY: Just a thought that comes to mind since we've now decided that we want the raw data to be available publicly, if we include in the document a statement that not only will it be available, but will be made available on an ongoing basis no later than x amount of time as it was collected, like 24 hours or 36. But no reports, just the raw data would be available. Would that be helpful, Robert?

Okay, so why don't we look at getting something of that nature in here so it is stated that it's going to be available in some timeframe. Just thinking on the fly here, is 24 hours a reasonable timeframe to say it would be available from when it's collected? From what I've seen and looked at various systems over time, if you're collecting it and you're eventually going to make it public putting a 24 hour clock on it seems reasonable. Anybody disagree with that? Or think it would make it a lot more expensive to build or anything?

UNKNOWN SPEAKER: It might make it hard for them to filter out some of the bad data. I would guess that some of that work with something going offline or having connectivity problems might be done by hand.

---

RUSS MUNDY: So, what would be a more reasonable timeframe?

UNKNOWN SPEAKER: 72 hours, give them a day or two to clean up things.

RUSS MUNDY: 72 hours after collection, would that get the data available to folks soon enough? Brad?

BRAD VERD: This comment is not necessarily about the timing of it. This data needs to be reviewed and whoever is collecting it, there is GDPR concerns, are you a data holder, there's all these things now. So I'm not sure how prescriptive we can be around this right now. I think we can make a recommendation that we want this done, but there are laws now that need to be followed.

RUSS MUNDY: Yes, Jeff.

JEFF OSBORN: I'm wondering whether we shouldn't defy Paul and consider 'should', we should get it within a day. Because I had a real life example yesterday where Howard pointed out a vantage point which I had not been using and showing an outage of a sort, it was a route leak that we hadn't seen. So the idea that we're going to be measured on 30 days of something and don't have early warning that our problem is accruing I

---

think is going to be problematic. So if we're putting in a request or a preference, I'd sure like to get it within 24 hours being collected so that we can do something about it.

PAUL HOFFMAN:

I'm going to agree with Brad here. I think 24 hours is too aggressive for situations where data has to be cleaned, which we've agreed that it has, either for vantage points going bad, especially if they're periodically bad, or where the collector, the one who is going to be doing the monthly report, looks at it and says, oh, I can see an anomaly here that wasn't tripped by a trip wire, but is clearly wrong on something, and wants to do something about it.

We already agreed in Section 4.8 that that collector has the capability of saying I'm not going to use this negative data for some reason that I will describe. Putting that data out earlier than someone might have had a chance to look at it I think will lead to more confusion. So I think the 72 hours, I think that's a reasonable amount. And again, if you are really concerned, you can set up your own collectors near where the vantage points are.

RUSS MUNDY:

So, we have several views here and there clearly are laws impacting, potentially impacting, we don't know yet, since we don't really know what the data is, and would a compromise be to essentially say that the raw data should be available within a week of collection?

---

BRAD VERD: Could you not specify a time and say it would be our goal that the raw data should be made available as soon as possible per regional whatever the rules and laws are of the governing body. May even state that we talked about trying to put a time on this, but realize that there are outside forces that will be put upon whoever the data collector is.

DUANE WESSELS: Okay, I've been taking notes. Alright, next one. So, expectations on data sharing. We need to add some text to the document that the official implementation of the metrics will provide the raw data to anyone in the interest of transparency, so we can add this new recommendation here, the reason being that we want any third party to be able to verify the metric calculations. My question to the group, again, since this is something we talked about a long time ago with gaming the system, by publishing the raw data, we're kind of saying that we're less concerned about that now, I feel.

JEFF OSBORN: I think that you will get some responding on the mailing list when this goes out.

DUANE WESSELS: I just want to acknowledge that this is a change and what the previous rationale was. Alright, next. So this applies to what we're calling the reporting. We've agreed that for the RSOs the reporting would have just a pass/fail indication, but for the RSS metrics we would report the actual numbers. And I have a note here, to provide rationale for this

---

decision, that's to be written, I guess. I don't have that written down here. I do think we need to come up with a good rationale for why that's going to be the case. Someone would ask why is it okay for one and not the other. Liman?

LARS-JOHAN LIMAN: Liman here, I think that one of the reasons is that some of the RSS metrics are complex combinations of the RSO metrics and just having a pass/fail doesn't give enough information to the reader.

DUANE WESSELS: Doesn't give enough information to what?

LARS-JOHAN LIMAN: To the reader, it doesn't convey the correct message.

DUANE WESSELS: I see, okay. Thank you. Alright, let's go to the next one. So, regarding the examples in the end of the document, we've agreed to keep those in the document but we need to be more explicit that the actual implementation, the output will probably look different and this is just sort of guidance for certain readers to help them understand what the output of the metrics is going to be like, what we expect it to be like. But the actual structure and format may be different and it's going to be implementation dependent. Alright, next.



---

Now we get into some of the specific metrics. So for RSO availability there were no changes to the formula or the method. We've settled on a threshold value of 96% here. We determined it business using this [inaudible] parallel availability formula with  $K = 8$  and overall ability of 5.9. We will add some text to the document that says that in the process of making these measurements, we are also measuring things that are out of control of the RSO such as the network components between the vantage point and the servers, and that's one of the reasons why it's not necessarily higher. Okay.

So, for RSO latency, again, the overall method is good, the thresholds we settled are on 250 and 500 msec for UDP and TCP. The rationale here is that this is approximately twice the earth's circumference at the speed of light, which if my math is correct, that's 267 msec, so we can say we rounded it down a little bit. Does that match everyone's recollection of this? Or did we just pull the numbers out of the air?

RUSS MUNDY:

I think it would be cool to have a reference to the speed of light in our document, yes.

UNKNOWN SPEAKER:

It's actually a very elegant little formulation.

DUANE WESSELS:

And then again note that TCP is twice UDP because of connection setup, which think is very generous.

UNKNOWN SPEAKER: Why TCP is twice UDP?

DUANE WESSELS: Because in the worst case, for TCP connection there is a connection setup roundtrip. The size of the data is not a big contribution to the latency. It's one roundtrip for the setup and one roundtrip for the query response. In TCP there is a tear down component, but we're not counting that into calculation of the latency.

UNKNOWN SPEAKER: I think that the first time they try to use UDP, if there is [inaudible] in UDP, it will be transferred to TCB. I think it will have more than twice.

DUANE WESSELS: That's true for a resolving that's doing things, but in this metrics speculation, we don't say that if you get a truncated response, so these are measured separately, there would be a UDP query and then a separate TCP query, not related to whether the TCP was set or not. So there will always be both in every interval. Okay, next.

For RSO correctness, there's a lot of changes here. We've agreed to combine these, the current two metrics into a single one, so this part of the document will change a lot. We need to think about whether the DNSSEC okay bit should be set on all the correctness queries or only on some of them. We need to update the rules for when to expect signatures, or when not to expect signatures, and so on.

There's a list here that says three types of queries, and that was based on the conversation that we had yesterday and then I know Paul threw me a curveball this morning that said he's changed his mind on this a little bit, so this may continue to evolve. But the idea is to have a mixture of query types to cover all the common cases of root server responses, as well as to generate queries that cover the gaps in the zone and try to find places where TLDs have been inserted that they don't belong. Paul?

PAUL HOFFMAN:

One change from part of the document that isn't listed here is did we agree that the correctness either matching or DNSSEC validation will not be done on the vantage points but will always be done on the collector of the data? Because right now the document has DNSSEC correctness being done on the vantage points, but matching being done on the overall collector.

DUANE WESSELS:

Yeah, I think it's a detail we didn't really talk a lot about, but my feeling is that this will all be done on the collector now, because I don't want the vantage points to have to maintain the zone file archive, I think that's too burdensome, and I think you agree. So I think this will all be done on the central collection system from now on. So, for correctness, we agreed that the threshold is 100% and we will add text to acknowledge, to be more clear, not that the thresholds can never be perfect, but what can be perfect is the system will never be able to find

---

all the cases in which a determined root server operator could falsify responses.

RUSS MUNDY:

I think maybe a better descriptive way of saying this threshold needs to be 100% is that every response observed by the vantage points will always be 100% of what it's supposed to be either from a correctness matching or from a DNSSEC perspective, rather than trying to say specifically detecting things, changes that occurred. If any change occurs that is include in a query set, it will be identified as a failure. Because we can't guarantee from this monitoring system that we could find any insertion.

DUANE WESSELS:

Right, that's what the last bullet is meant to acknowledge, but the output of this metric is a number, a percentage, and our threshold for pass/fail on that percentage is 100%, right?

RUSS MUNDY:

Yes.

DUANE WESSELS:

Okay, so you're right, we will acknowledge that we can never be perfect in the coverage in identifying all those cases. It will always be possible to cheat, unfortunately. Okay, next.

---

Pretty straightforward for RSO publication latency. There was no need to change the method, right? We didn't do into much detail about it because it's very complicated. But at a high level there is no need to change it. The threshold we set to one hour and the rationale for this is that it's twice the SOA retry value. Keeping in mind again this is a median metric. Alright, next.

Now we get into the good stuff. The RSS availability.

PAUL HOFFMAN: Pointing out from what Suzanne said earlier about having physics, now we have physics and math, this is all good.

DUANE WESSELS: This is designed to reflect what's on the white board over here. So we have agreement on the value of  $K$ , so for  $N=13$ ,  $K=8$ . In each measurement interval  $T$  and for each vantage point  $V$ , we calculate the number of RSOs that responded to an availability query. If that little value of  $K$  happens to be greater than 8, then we set it equal to  $K$ , so that's the maximum value in any measurement interval. And then the calculated daily RSS availability is the sum of all the little  $k$  divided by the sum of all values of big  $K$  out of the maximum possible value. Does that make sense to everyone? That matches what we talked about here. Robert?

ROBERT STORY: I just think we need to make sure that we define availability here, since it's now different than what RSO availability is. So we're saying that it's

---

available if we get 8 or more answers, and the math is good for explaining how we came at that and how we're calculating it, but we need to say that this is not the normal definition of availability.

DUANE WESSELS: Well, in my mind, this is the definition of RSS availability, this is how we're defining it.

ROBERT STORY: I think we should also find a way to maybe put it in simpler language than making somebody look at that formula and go back to their college math books.

DUANE WESSELS: Okay, certainly we can do our best there and happy to have your review of what this section ends up looking like. Yeah, it won't be as terse as this, this is just to sort of fit on the slide. As we talked about yesterday, we're going to include some examples which actually I have on the next slide, so we'll go through those, as well. Did you hand go up, Ryan?

RUSS MUNDY: I wanted to make sure because Robert brought it up here, we have agreed to not change of this metric, we went back and forth yesterday, there were several proposals and several discussions about changing the name of this, but we are still concluding it is RSS availability.

---

DUANE WESSELS:

Okay, so since this is a tricky one, on the next slide I've got some examples and we'll put these or even more like this into the document. So in the first example, this is the best that you can have that's below 100%. So, in one 5-minute interval, 1 vantage point could only reach 7 RSOs and in all the other intervals, all vantage points could reach at least 8 RSOs. The calculated RSS availability is 99.99992%.

Next example, there is a 24-hour attack that takes out all the RSOs entirely, one day out of a 30-day month takes out all the RSOs, the calculated RSS availability is 96%. A month-long attack, it takes out 6 RSOs entirely. So instead of reaching 8 as the minimum, the vantage points can only reach 7, all vantage points for the whole month can only reach 7 RSOs, the availability is 87%. A month-long attack takes out 5 RSOs entirely, all the vantage points can reach 8 still, the availability is 100%. Make sense everyone?

JEFF OSBORN:

Just kind of being the devil's advocate, I know we walked through all of this mathematically, but doesn't the disparity in terms of the size and numbers of instances of the RSOs mean that putting that many significant digits on something, it's a little bit of a fallacy. If the 13 RSOs were identical, this makes perfect sense. To the degree that there is a disparity in them, it's harder to buy. Now did I just miss some piece of math where you guys all decided it either makes sense or it's a detail? Because it seems like the first thing somebody is going to bring up is like which 5, which 6?

---

DUANE WESSELS: That is a simplifying assumption that we have to make here, that all RSOs are identical. If they're not identical, it gets really, really complicated.

SUZANNE WOOLF: This is Suzanne. I think part of the difficulty, I've been sitting here thinking there are going to be some challenges of explaining all this, and that's a separate problem from what we're solving now to do this work. Because explaining any of this in terms of user observable behavior is going to be difficult to impossible and that really has nothing to do with the metrics we're trying to do. But I think you're right, there are going to be some challenges with explaining it, and that's actually a separate problem to what we're doing now.

UNKNOWN SPEAKER: So, two things, one is before all the math, we said that we came up with 8, and I think we may also want to go back and explain how we got to that rationale, which was Paul Vixie's description of wanting to get the next request. But looking at what we're saying now, where we're talking about putting out availability numbers for the month, I think Paul's rationale is not valid. Because he was talking about wanting to have 8 so that when these RSOs go down, the next query is going to have a good probability of getting an answer. So we're talking about the correction of a resolver which is going to happen within minutes. It's eventually going to figure out what 5, 4, 3, 2, 1 are going to happen. That's going to happen on the order of minutes, and we're reporting availability based on a month.



---

So even if you get down and do the math, not this math, but the math that I think Andrew was talking about, where adding up the number of queries sent, the number of RSOs responding, for the typical, the same math we're doing for RSO availability, whether or not answers were reached, even if you get down to RSO within a couple hours everybody is going to be talking to that RSO. So on the scale of availability for a month, I don't know that the rationale of saying we want to have at least 8 so that the next very query that a resolver sends is going to have a high 2/3 probability of getting an answer. The scale of the rationale is way off.

DUANE WESSELS:

I guess I don't see it that way, but I appreciate this perspective. I think based on something I heard you say yesterday, in your mind the system is available if there is response from at least 1 server, we could certainly take that approach, but I think the risk there is that it doesn't capture the situation when you've got half of the servers running and that would be a concerning situation and you would want to alert on that as a metric, as a threshold.

UNKNOWN SPEAKER:

Right, and I agree. And that was one of the things when we had the discussion yesterday where I talked about the possibility of having two metrics. One, is the system available, and two, is the system healthy. So even if you've got 1 RSO eventually everybody is going to converge on it and the system that's available might not be responding as quickly as we'd like under the load, but it's available, but it's not healthy.

---

Where I could definitely say the threshold for the health of the system would alert much quicker than getting down to just 1, even getting down to 10 I think is troublesome, worrying for a health threshold. But from the very beginning of the work party like I said and several others on the mailings have said, if you're getting answers the system is available. That's the whole point of having all the redundancy. So just getting answers from one place is not the best place to be, but the system is working as designed.

BRAD VERD:

So, yeah, if you get an answer the system is available, but it is not reporting availability of 100%. If you define availability as 8, which we've done in here, and you're down to 1, you are no longer at 100%. The system is available, it tells you a percentage of what is available. The whole reason we did this is because you didn't like if you went below 8, the system was not available. So we did this to accommodate that, which was reasonable, I think, but I think it is completely unreasonable, as availability is not something that's new in this room, the thought process around availability is well known throughout the world and it is not all or nothing. Right now you're saying that if there is an answer, it's 100% available, and it's not.

KARL REUSS:

I think it comes down to the definition of availability and what we're saying is this is the percentage of time that you will get a timely answer to a single query. And if 4 are down, we've still got 8, and it's still 100% chance of getting a timely query, you lose the first one, you've got a

---

good probability of getting the second one. So it's just the amount of time you've got a good chance of getting a quick answer. You can try again, and you'll eventually get your answer.

PAUL HOFFMAN:

I was going to pretty much say what Karl said. I wasn't going to say timely, but reliable. I was one of the people who was picky about saying let's define availability and as we got into the discussion I realized, no, let's not. That is not something, if we define it based on the way we think, the general public won't get it, if we define it the way the general public get it, then these numbers will be meaningless and we won't be able to actually help ourselves. So I'm actually in favor of not trying to say timely or reliable or whatever, and just say it's availability and describe how we got it.

Remember, all of these numbers will be out there, so someone can say I don't like your definition of availability, take our numbers, and create their own report. And if somehow that gets consensus, we might change, or add it, or something like that. But it's not like we will have any monopoly on the definition of any of these thresholds, people can have their own thresholds, including RSO thresholds. They might set up a system where there is going to be a lot more red because they like red better.

DUANE WESSELS:

So, I'm just going to have to disagree with Brad on his definition of availability. For the RSOs, we're saying if we get an answer from an instance, we're using the address, that it's available. It doesn't matter if

---

that address is a load balancing box that has four boxes behind it, we're not saying we need an answer from every box or it's not 100% available. If we get an answer when we ask the question, it's available.

BRAD VERD: How is the RSS defined in our lexicon? The availability metric here is for the RSS, so it's the 13 identities, it's not 1 identity, this 13.

DUANE WESSELS: Then why are we saying we need an answer from every one?

BRAD VERD: We're not. This is the availability of the root server system, and we're not we saying we need an answer from every one, we as a group said we need an answer from 8.  $K=8$ .

DUANE WESSELS: So we have a different definition of availability then to general public, I think, so that needs to be defined. And again I think it's like Russ said, if we change the name of the metric from availability to something else, it wouldn't be an issue.

SUZANNE WOOLF: I think that's actually back to what I tried to say before about trying to describe, all this discussion has led to these metrics as being useful in terms of what we're trying to say to each other and about each other to the world, but the problem of what does it look like to a user,

---

translating any of this into what users see or don't see is going to be difficult and it's a separate problem than what we're doing here. This one is going to be particularly difficult because yeah, it's going to sound weird to people to say 87% of what, either I've noticed no issue, so how is there an issue, or I can't get what I need, how can you say that the system is mostly available. But turning any of this into user visible experience is going to be really hard.

LARS-JOHAN LIMAN:

Liman here. Is there a missing bullet up at the top of this document saying something along the lines of the DNS as a whole has so many moving parts and some many time constants, and so many axis modes, that the entire system is not fully predictable. And what we tried to do here is to give some values that can help us assess the health of the system. But that unpredictability thing, it would be good if we could get that message to the reader.

DUANE WESSELS:

Yeah, that seems like a very good thing to add to the document, I agree. But I don't know if it gets us past this thing that we're stuck on right now. Anyone else want to chime in?

PAUL HOFFMAN:

My own interpretation here. We're defining the RSS availability at capacity, that  $K=8$ . I understand if you get a response to the RSS the user will say it's available. I get that. I think it's good what Liman said, we're not trying to describe the user experienced because there are so

---

many moving parts between us. This is what we're defining for our own measurement of our own system which is one piece of the entire user experience. So we're redefining availability specifically this way [inaudible].

DUANE WESSELS:

Yeah. One thing I struggle with a little bit with, what you're advocating for, Robert, is that -- it does make some assumptions about the way recursive name servers work. I mean, we do have some general understanding of that but it's not something that's -- there's no RFCs around how they have to retry it and things like that. And after how much time, and when, and things like that. Matt, you had a point?

MATT LARSON:

Yeah. I think it's valid to extrapolate from K=8 to month-long availability. And maybe if we thought of availability in terms of usability, the system is available if it's usable. We all, based on our experience as well as some gut feel, decided that if we got below K=8 there would be at least some cases where recursive resolvers would time out. So that would effectively make the system unusable by a significant population. So, I don't have any trouble seeing that as unavailable.

I think by some strict definition of "I will eventually get a response", if that's your definition of availability, then I see where you're coming from. But that's sort of a pathological availability because it's not useful. But if we think of availability in terms of usefulness, then I think

---

this all drops into place and then we have something that we can get our minds around.

BRAD VERD:

Adding to what Matt just said, I looked up the definition of availability and it's, "Sustain a system or capability at an agreed level of readiness," or you can say readiness is usability.

RUSS MUNDY:

One of the things we discussed specifically about the RSS availability is if we used a different name from the RSO to the RSS. Based on Matt's suggestion, I think just a quick look at what's currently there for the RSO metric, we could put usable availability in both the RSO and the RSS metric. And I think it would make sense -- is something that would be a viable solution here. Robert, you had your hand up, so go ahead.

ROBERT STORY:

I think we need to move forward, but I'll say two more things and then hold my peace until discussions happen on the mailing list. One, talking about availability. Again, the definition that Brad just gave is at a defined level. So what we haven't sufficiently explained is at the defined level. Somebody else said something about other people can measure. Well, other people can measure at eight RSOs responding. We're saying that it's 100%, but if you actually measure the number of queries sent to all the RSOs and the number of answers you're getting at 66%. So, other people looking at the definition of availability are

---

going to have a different interpretation of that. So we need to be very clear about defining what it is.

BRAD VERD:

I agree with you. I feel like we've defined it at eight and as you just said, somebody else outside of here could define it at 13. But we've defined it at eight or one, but we've defined it at eight.

DUANE WESSELS:

Right. But I think maybe what Robert's getting at is also, the part that's missing, we don't have the -- because it's all new -- but we need to say not only at eight, but with some expectation on timeouts and retries and that sort of thing. Kind of getting back to the two thirds probability thing.

ROBERT STORY:

Yeah. And that was actually my first point that I wanted to make that I forgot. I think early on we had discussions about whether or not we cared about how resolvers behaved. And I thought we said that we didn't. And now we're saying that we do. And I think that if we have a resolver work party, I will be very surprised if there are resolvers out there that do not very rapidly converge towards a number of name servers that's lower than eight, for being able to function and get an answer every time.



---

UNKNOWN SPEAKER: Could the test probes not actually include one or more common recursive resolver implementations and then actually measure those for real? You'd probably want to set the minimum TTL on any records they receive so that they don't cache them unnecessarily. It shouldn't be impossible to achieve that.

DUANE WESSELS: That's right. That is something that we did talk about quite a bit earlier on in this work party, but we took that out in the interest of simplifying and focusing our work. So at this point, we don't have anything in the document that says the vantage point should run recursive name servers and measure that way.

UNKNOWN SPEAKER: I was thinking, in addition to the authoritative tests.

DUANE WESSELS: Right. We were considering in addition, if I remember correctly. Or it was an alternative way of measuring it. So, I don't know. If that's the direction the work party wants to go, we can add that back in. It feels to me like it's going to delay progress on this quite a bit. That's taking a big step again. We can save it for version two or consider it for version two.

UNKNOWN SPEAKER: I see where you're coming from. I can't see how else we can replicate this user experience of is the system down or not.

DUANE WESSELS: Well, again, the goal of the work party is not to replicate the user experience, although I feel like in order to get where we need to get, we have to step a little bit in that direction. We have to make certain assumptions about the way the users of the system function. But in general, we're not trying to replicate the user experience in these metrics. These are designed to advise the operators or the oversight bodies of the operators, whether or not the system is functioning at the desired levels. Paul? Daniel?

DANIEL MIGAULT: It's just one general comment about the choice of  $K=8$ . What I'd like to raise is that when we choose eight, we have a rationale for choosing eight. But it's a very, very high bar. I think that's something we should insist on. When we say we choose  $K$  which we're putting the bar pretty high to say the system is available. I understand the rationale, but guess what? If the application doesn't have it in the first round or the second round, it's going to refresh. So, I think we are [inaudible] the problem that may not exist. It should be considered just as the rationale, and we just have to mention that the bar is pretty high. That's, I think, what should be in the text.

DUANE WESSELS: Okay. Thanks. All right. Should we move on? As Robert said, when we go back and do the work and update the document, take a careful look at this section because it's obviously contentious, and we want to get it right. So make sure you're happy with what is written in the document.

---

---

So the next is RSS response latency. This also is changed from the current method in the document. What we settled on here was -- again, this uses the value of K and its measurement. So at each measurement interval, each vantage point will -- well, it'll analyze the latency results and it'll find the K best response latencies. So it'll throw out the ones that are above that. And then for the aggregation interval, it calculates the median of that subset of all those latencies and all those best latencies.

FRED BAKER:

Aggregation interval here is also a month?

DUANE WESSELS:

A month, yeah. We didn't spend a lot of time talking about threshold values on this, but at the end we heard people throw out as a proposal, 150 milliseconds for UDP and 300 milliseconds for TCP. We need to come up with a better rationale than this, other than the fact that they're just lower. So, maybe we'll come up with some math behind that. Maybe there are one third lower or something like that. But that's the idea. Comments?

Okay. Next is RSS correctness. Again. We'll modify these to use the new RSL correctness definitions and measurements. The aggregated metric would be a simple fraction of number of correct divided by total number of responses. The threshold is 100%. The rationale is that we must always be correct. It must be perfect.

---

Okay. Next, RSS publication latency. This does not currently appear in the document. It was taken out before we agreed to put it back in to maintain symmetry and to have it in there even if it's essentially the same as in the RSL case. So the threshold is the same, one hour. But we'll keep it in so that its absence is not glaring. Robert.

ROBERT STORY:

I would say the same rationale for having the response latency lower for the RSO applies here. Again, we're talking about a median, so the out layers are going to get tossed out. So I would suggest maybe that the SOA refresh timer, without multiplying 30 minutes, would be appropriate.

DUANE WESSELS:

Okay. I like that. Anyone else have opinions about that proposal? Okay, I've noted that down. Thanks. All right, next slide. And this is actually the last slide.

So, we've talked about some future work items and we'll add a section to the document to capture those. Some of the future work items are directly related to the metrics work. One of the things we talked about was adding RSO self reporting, possibly in the future. That has the advantage that you can cover more components, probably, and you eliminate uncertainty about components that are not under the operator's control. We've already talked a number of times about investigating ways to improve vantage point distribution, especially as they grow in number. We talked a little bit about a reference data set.

My thought here is that we'll have a GitHub repository with a reference data set that people can use to test their implementations against or their aggregation methods against. One of the ideas here is that we want to make sure that everyone's interpreting the document in the same way and that everyone gets the same answers out of the reference data set. In my mind, the reference data set would include some cases where there would be failures. So we would maybe force the data a little bit to have cases where it's not always going to meet the thresholds.

Some other things we talked about separate from the metrics work party would be the financial aspects of this and its relationship to the PMMF. Uh, Jeff, your name is mentioned here specifically. This is something that you're passionate about. And we also -- what was our question?

JEFF OSBORN:

It's really the next one. We need to probably put SAPF as well as PMMF because it's SAPF that defined in 37 as the one that says how many we need. So we probably got to point to both of them.

DUANE WESSELS:

Okay. Yeah, thanks. So, there may be a future document that advises those functions or whatever they become, on how to interpret an act on the data from this. And as we can see, that'll be very important because even amongst ourselves we're having some difficulty understanding how they should be interpreted and what they mean. So that would be very important.

---

---

So that's the end of the summary slide that I have. If anyone has thoughts on something that was missed or that we need to capture, please feel free to speak up -- or any other comments about what we talked about today.

PAUL HOFFMAN:

We're going to have discussions on the mailing list. I'm trying to be polite here, but previous discussions on the mailing lists have gone almost nowhere. And yet I think that we have a bunch of topics where changes that we all agreed need to be made, but we're not exactly sure how we're coming up. So hopefully, people are more active on the mailing list and that we keep subject titles to be relevant. It seems likely there'll be four or five threads running at once. Once you start doing things, and we see a bunch of them, then someone's going to say, "I want to talk about this one or this one." And if you just say "metrics document", that's not going to help.

I have felt underwhelmed by the use of the mailing list before. And you know, most of us also hang out in IETF where it's easy for a mailing list to be useful and it's really easy for it to be not useful in either direction - way too much conversation or you aren't sure if silence is consent or boredom. Because I know I'm going to be contributing on a bunch of these, so I'm doing this as in when I say stuff I would like to see responses, not just someone thinking, "Oh yeah, that's fine. I don't need to say anything." Because it's probably not fine, as we've discovered already with some of the assumptions that we are reversing here.

**RUSS MUNDY:** In support of what Paul was just saying, I think that one of the things we have got pretty strong agreement on is our structure of the document by section and title and names of the metrics. And even if the names of the metric change, we have a current name. So perhaps a informal standard that we could adopt as part of how we discuss these things in the future is in the subject too, for under discussion. Try to relate it to a particular metric or a particular section -- like 4.8 for the exceptional events or something -- to help keep focused on the topic. Because I think there still will be remaining things and as we get changes put into the document, people are welcome to make changes to the text of the document. Please do so. And we can discuss the extended discussions on the mail list. But making changes to the document is a good thing, not a bad thing.

**DUANE WESSELS:** I think that's it for us. I hear activity out in the hall, but we weren't scheduled to break until 12:15, right? Is this really the last thing on -- Did you want to do a workshop wrap up, Fred, after lunch?

**FRED BAKER:** I don't see a need for the wrap up to wait until after lunch to happen. Frankly, I think we've made a lot of progress in the last three days. You guys have a lot of writing to do. And I appreciate that. Thank you very much for doing that. As that goes forward, I'm going to suggest that we collectively watch those Google docs. There'll be further things to comment on. And with any luck, we'll be able to finalize that by the end

---

of the year, which I think is a good thing. I'm not sure what else there is to actually comment on. Brad, do you?

BRAD VERD: Future of workshops?

FRED BAKER: Future of workshops. Well, the way I understand it, ICANN has allocated money for us to have two workshops like this next year. As you mentioned, I could imagine the GWG using one or both. I don't know that I've got a topic right now that calls for us to have this workshop, which was the point behind my comment yesterday. If we hadn't made the progress that we made this week, I would have been seriously thinking about a workshop in the spring to accomplish this -- "Okay, we've accomplished this." So I would expect that the next workshop is probably GWG. We'll see. Changes always happen, but that's what I see right now. Brad, do you have --?

BRAD VERD: Yeah. People keep asking about workshops as if we want to cancel them. I'm hesitant to say no to any workshop because it's when we do most of our work, face-to-face, short of the ICANN meetings. So, I'm hesitant to say no workshop because I believe that we'll have something on deck that needs our attention. But that's my personal view.

LARS-JOHAN LIMAN: That's what I hope to hear.



FRED BAKER: You hope to hear that we have nothing on deck or that we have workshops?

LARS-JOHAN LIMAN: That you are not cancelling workshops, because we will have something to talk about.

FRED BAKER: Oh, I'm not canceling workshops. I'm just saying I don't know what the topic is going to be. [AUDIO BREAK]

So, next on our agenda, we're scheduled for lunch in half an hour. I would suggest that we adjourn. People are welcome to have lunch. Brad and I and staff have a meeting, and then we're done.

LARS-JOHAN LIMAN: I just wanted to extend my warm thank you to Verizon for hosting this and to ICANN for arranging all of this. So, thank you very much.

BRAD VERD: We're happy to host. We're happy to help.

**[END OF TRANSCRIPTION]**