
DUANE WESSELS: Good morning, good afternoon, good evening. This is the RSS Metrics call held on the 22nd of July 2019 at 12:00 UTC.

In the Zoom room today we have Abdulkarim Oloyede, Dessalegn Yehuala, Duane Wessels, Jaap Akkerhuis, Paul Hoffman, Russ Mundy and Shinta Sato.

From staff, we have Andrew McConachie and myself, Ozan Sahin. I believe we have 11 participants in the room, so if you can go ahead, go around the table, and if you can identify yourselves, I would appreciate that.

Also, let me remind that if anyone steps in the room, could I please ask this person to identify himself or herself? Thank you.

DUANE WESSELS: I'll start. This is Duane Wessels, co-chairing this meeting.

MAURICIO VERGARA: Mauricio Vergara.

UNIDENTIFIED MALE: [inaudible] ISI.

UNIDENTIFIED MALE: [inaudible].

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.

UNIDENTIFIED MALE: Ozan from ICANN.

FRED BAKER: Fred Baker, ISC.

SHINTA SATO: Shinta Sato, JPRS.

KARL REUSS: Karl Reuss, University of Maryland.

KEN RENARD: Ken Renard, ARL.

BRAD VERD: Bard Verd, Verisign.

RYAN STEPHENSON: Ryan Stephenson, DISA.

UNIDENTIFIED MALE: [inaudible].

DUANE WESSELS: Did you get everybody, Ozan?

OZAN SAHIN: Yes. Thank you. Over to you, Duane.

DUANE WESSELS: Okay. So thanks, everyone, for making the time to come to this meeting. We've been meeting about every couple of weeks and having some good discussions.

Today, there's a few things that I had in mind to talk about, and this is a little bit different than maybe some of the things we've been – well, in the past, we've been focused a lot on the nature of the metrics, like round trip times and aggregation, stuff like that. But we've gotten to a point where we're a little bit stuck, and I think we need to talk about some broader issues.

One of the things I want to talk about today, we've already put out a message on the list, which is sort of the purpose of the metrics. There's the proposal to resolve that.

And then following that, we'd like to spend some time talking about some issues that are a little bit related in my mind. One of them is who will operate the probe infrastructure. Will that be an independent third party, will it be self-operated by the root server operators?

Little bit related to that is some sense of coverage, how we want to sort of – [inaudible] we want to cover the root server system as a whole, we ought to maybe target specific [inaudible] or treat it as a cloud and let Anycast [inaudible]?

And then also related to those two things are locations, where might probes be located. So I don't necessarily expect decisions on all these things today, but we've had discussions. In the room here, we have some folks who are root server operators and have maybe not participated in the work party so far, so I thought I'd spend just a little bit of time with background, talking about where we are today with the work party.

As I said before, we have a document that's in progress, and there's a fairly good amount of material in there about the types of metrics that we're trying to achieve here. Really, there's four types there. Availability, latency, correctness, and staleness or freshness depending on your point of view.

Availability is pretty straight forward. That's how often you get a response to a query. Latency on one hand seems sort of straight forward. It's measuring round trip times, but it's really complicated by some of these issues that we'll talk about today and it is complicated by locations and that sort of thing.

Correctness, in the document we've sort of focused correctness as the measurement that uses DNSSEC validation, so the idea is that you do a query [with the submission of the response] and you validate it, and if it's good, then you say it's correct.

There's some metrics in there that tried to sort of spread out, or in some cases you're doing queries for names that exist and some you're doing for [inaudible] try to cover all types of responses and [inaudible].

One thing that we talked about with regard to correctness is that location matters, and if you're doing measurements over infrastructure that's out of control of an operator, then you might expect problems, you might expect that in some places there is interference with your packets, and that would affect your measurement results.

And to the extent we're talking about service level expectations, that's maybe not totally fair, it's not such a great idea to hold an operator responsible for things that happen on infrastructure that they can't control. So that's sort of an outstanding issue.

The last one was freshness, staleness. Here we can focus on just like looking at the zone serial number and seeing how [inaudible] up to date. This is relatively straight forward, except in my mind, the complication is you need either to know sort of the absolute truth like when [the zone] should be propagated to all the servers, or you need to do some kind of collective measurement and say, "Oh well, 50% of the things we measure now have this so therefore that's the most recent one." So that's just a little complication there.

Any questions or clarifications on the background and where we are so far? And Russ, please feel free to chime in too if I miss anything.

RUSS MUNDY: Sounds good. Thanks.

DUANE WESSELS: Okay.

UNIDENTIFIED MALE: [inaudible] the right answer, but [inaudible]?

DUANE WESSELS: In some of the proposed metrics – I don't think correctness is one of them, but the availability and latency [cases,] we are proposing to do those over all the permutations [of transports] so you'd have a different value for [inaudible] v4 availability and UDP/TCP – or v6 availability.

For correctness, we're not really suggesting to do that, and it's not designed to look at fragmentation or anything like that. It's really designed to look at, did I get a response that validates the DNSSEC signature?

FRED BAKER: I've been thinking lately about – wondering if we need to test the correctness of the zone as in what's being delivered. All of the records that have been produced by IANA, and only the records that are produced by IANA, DNSSEC doesn't tell us that.

But I've read an article the other day, [said that –] and I don't know this to be true, I only know the article said it – that text records [inaudible] floated to [DOH] servers and being used to control botnets. I wouldn't want that [inaudible].

So, do we have a way in there to know that we've got exactly the set of records that we should?

DUANE WESSELS: I don't think so.

UNIDENTIFIED MALE: Well, with DNSSEC or [inaudible]

FRED BAKER: DNSSEC tells us about individual records, [inaudible].

UNIDENTIFIED MALE: [The article describes talking about] would be for things that are not TLDs that would be [inaudible] domain names.

UNIDENTIFIED MALE: [inaudible] botnets and [DOH] servers were doing it to ICANN-controlled zones, or like their own domain [inaudible]?

FRED BAKER: DOH in terms of Google and Amazon and such download zones from various [inaudible] places. One of them is the IANA. And I think they also get [inaudible].

And then the server has whatever database it has, and when it's asked a question, it gives an answer. No, they're not controlled by ICANN, they're just using data in terms of the IANA zone. [inaudible].

DUANE WESSELS: [inaudible] but you're talking about [inaudible]. Basically [inaudible] being inserted into [inaudible].

FRED BAKER: Yeah. If you're delivering something and you think that it's part of the IANA zone, I'd like to believe that IANA agreed with [inaudible].

DUANE WESSELS: Right. So we don't have anything exactly like that right now. That's something we can consider – you could do – [since it's a DNSSEC zone,] you could walk the entire zone. You could do that. That would be [some assurance] I think as to what you're asking. There's also this [Internet] draft I'm working on called [message] digest for zones, which sort of touches on this.

Really, if you wanted to use that, you would probably have to do a full zone transfer [inaudible].

FRED BAKER: Okay.

DUANE WESSELS: But at this point, that's not really something [inaudible].

FRED BAKER: So is that a hash of the zone or something like that?

DUANE WESSELS: It's a hash of the on-the-wire content of the zone.

UNIDENTIFIED MALE: So one thing related to what Fred just said that I have put comment in, if we only do DNSSEC validation. We are not checking any [inaudible] records. I don't know if we consider that important or not. And what I suggested was if whoever's testing actually has just pulled a full zone from a place [they know,] they've also got the new records we could check [inaudible]. And then if they need to do DNSSEC validation. I don't know whether that's that important or whether people – like we consider that to be part of the SLE/SLA of you need to be doing also the additional records, or as Brad said, somehow, if we could possible check for every possible name we thought that might appear there, do a query and see if they get an answer. Don't know.

Certainly, both would work. DNSSEC validation works, so does just looking against a known good copy of the zone.

DUANE WESSELS: Right. Andrew, your hand is up.

ANDREW MCCONACHINE: Yeah. I just wanted to remind people to please state your name before speaking. We have a lot of remote participants.

UNIDENTIFIED MALE: [inaudible].

DUANE WESSELS: Okay. Thanks, Andrew. Russ.

RUSS MUNDY: Yeah. Thanks. It's a little hard to hear some of the people in the room, so if you could speak relatively loudly and not too fast, that would be real helpful to be able to follow that discussion. Thanks.

DUANE WESSELS: Okay.

FRED BAKER: Maybe shove the hockey puck a little bit closer to those folks. I don't know if that'll help.

DUANE WESSELS: Alright, so I want to move on to sort of the first item that I talked about in the agenda, which is the purpose of the metrics. There was a message that I sent out on behalf of myself and Russ and [Shepherd's] – I think it was on Friday – basically proposing that moving forward, let's look at the work party only on metrics as they relate to service levels.

So in previous meetings, we've had this discussion in – I feel like it's sort of – we fall into a little bit of a trap where we say, "Hey, we're doing these latency measurements. We could also use these latency measurements to do these other things, like tell us about the behavior

of resolvers or we can provide data to researchers to learn about the root server system and so on.”

But I think that that’s causing us to sort of get stuck. So given that the sort of charter or the statement work for this work party was kind of derived from the RSSAC-037 work and the idea that there may be service level expectations on operators in the future, let’s focus just on that.

So that’s the proposal [inaudible] work party leaders and chairs. If there's any objections or concerns about that, now would be the time to state them, otherwise I think we’ll move forward and sort of work from that perspective and start modifying the document to reflect that.

FRED BAKER: [inaudible] sent a note overnight on that topic.

UNIDENTIFIED MALE: [I hadn’t seen it.]

FRED BAKER: [inaudible].

UNIDENTIFIED MALE: [inaudible]. In my opinion [inaudible]

UNIDENTIFIED MALE: Should not depend on [inaudible]?

UNIDENTIFIED MALE: [inaudible] having some context for the service levels. And I was [inaudible] as you propose. Also doing measurements of something else besides the root servers to get a better idea of the general service level than specifically ... it takes a really long time to get to the root servers and it takes a really long time to get to TLD servers. [inaudible] it's the service level between the measurement point and the node that's [maybe a slow lane in a remote location as opposed to latency that's due to the issue with the root server. That's at least one issue that we thought about too.]

DUANE WESSELS: Okay, so does this really relate to the question of the purpose of the metrics though?

UNIDENTIFIED MALE: It's not the purpose, but if you're talking about the [inaudible] service levels based on the metrics, I think the context of verifying the service level is not – make sure we're looking at root server versus there's no connectivity [inaudible].

DUANE WESSELS: Okay. I can understand that. To me, that's still a separate issue from saying that the metrics should serve the purpose of – so I think it's compatible with saying the metrics are [inaudible] service level.

UNIDENTIFIED MALE: [inaudible].

DUANE WESSELS: Right. What I'm proposing or what I'm talking about is let's not [inaudible] can be used for other purposes [of measurements.] Let's focus just on the service level [stuff.]

Now, to say that the metrics should be relative to [– connected to the data as a whole,] that's fine, we just need to come up with a way to do that. There needs to be a concrete proposal for what kind of measurements you're going to do to compare against. And then we can have the debate about, oh, we should be doing round trip time to Google, to Microsoft, to whatever, then you have to have that argument about what's our basis, or whatever technique.

What's proposed so far is a little bit simpler, it's just round-trip time back and forth. And also, we still haven't gotten to the thresholds discussion yet, and there was – I think David made this point at the last meeting, [inaudible] it may be the case that we end up with very high thresholds that essentially everyone can meet, in which case there's maybe no reason to have this [inaudible] comparative metric time to the rest of the Internet. So I feel like that is a future discussion.

We have some new folks in the room. We've got Jeff Osborne, we've got Daniel Migault, and they're sitting farther away because the room is getting full, so they'll have to speak up extra loud when they want to talk.

UNIDENTIFIED MALE: [inaudible] looking, Brad, to your [inaudible]. The issue of thresholds and the issue of somebody paying RSOs in order to do some level of [thing] are the two outcomes of this that are sort of more than academic. They have real weight to them and real consequences.

So if we were to say [inaudible] I think we need to recognize that we're doing it and recognize that that's either saying this [inaudible] to go or a different group has [two issues] to try to solve that are similar to what this group is trying to solve. Because those [inaudible] Brad, [inaudible] artificial deadline of any sort?

BRAD VERD: [inaudible] I think it's time for this group to [inaudible] and talk about how to do this and how to do it right. Sorry, I'm just going to say it. And the fact that [we can't have discussion on thresholds is fascinating to me] and also disappointing at the same time. So I'm sitting on the sidelines waiting for the discussion. [inaudible]

UNIDENTIFIED MALE: Okay. I'll be the other jerks who feels the same way about the SLA component of this where I sort of volunteer but [but I feel a] sense of ownership in a good way about the [SLAs] where we're talking about and ICANN is taking seriously the idea of financially helping out with something in return for something that they feel is important. We kick that can down the road or we can address it.

I've seen things [inaudible] budgets of theirs, like this was something they're going to start [writing checks for] in 2020. If I've misread that,

somebody else correct me. But that means we kind of need to have a handle on what we're doing is something that could become an SLA.

And to Brad's understanding, it's disappointing. On the other hand, it's hard. Thresholds is a hard discussion. And the SLA part is a discussion we're not used to. [inaudible] Internet and [inaudible] paying customer. [inaudible] paying customer [inaudible].

BRAD VERD:

I said something [to Duane] weeks ago [inaudible] maybe there are two sets of metrics. It's a set of metrics that are the health of the system and what does it look like, what we've been talking about, and then there's another set for [the SLA purpose.] [inaudible] maybe [inaudible] set of threshold numbers. [Hold on, let me finish.] So the second set of SLA numbers or metrics, let's call them, maybe aren't monitored for those [inaudible] because you remove this whole thing of Internet latency introduced by latency I'm not in control of so you can't hold that against me.

So this is how most of the TLDs do it now. They self report SLA numbers [inaudible] in their contract. And the number starts when the packet enters their network and it ends when [inaudible].

UNIDENTIFIED MALE:

[Do they ever get audited?]

BRAD VERD:

Absolutely.

UNIDENTIFIED MALE: Well, then that's [inaudible].

DUANE WESSELS: Russ has his hand up. Go ahead, Russ.

RUSS MUNDY: Yeah. Thanks, Duane. One of the things I've been thinking about this after participating in the discussions and so forth is that as an RSSAC caucus work party and the RSSAC as a group, the things that we're really in a position to state with some degree of authority are service level expectations. The expectations, not the agreement, not the contracts, no financial. We don't have any money, we can't give anybody anything.

And in the interest of, as the co-chair trying to help us progress, this document, this is one of the reasons why I've been encouraging having the approach – I think it's similar to what Brad was just describing where this document and this work is explicitly focused on the service level expectation idea. If others who are an authority to contracts have money, penalize people if they don't do X, if they choose to make use of the service expectations that come out of this group, that's between them and those that they write the contract with.

So my thinking has been if we focus and try to keep our focus from a service level expectation basis for this work party and let whoever the right other folks are get into the writing of contracts and making of service level agreement kind of expectations, I think that might help us make more effective progress in our work party here. Thanks.

DUANE WESSELS: Alright. Thanks, Russ. [inaudible]. I just want to ask people to be quick, because there are other things we want to talk about, and I don't want to [inaudible] stuff. So the question before the group is, should this work party focus only on service level expectations or not?

DANIEL MIGAULT: We should focus on one set of numbers, and I would like to see those numbers used generally everywhere, because having different sets [inaudible] going to make the world very complicated. At the very least, [we should try to explain it to others.]

BRAD VERD: I think we should have one set of numbers. I think we shouldn't be focused on the research stuff to the question at hand. TO your comment, Russ, I kind of disagree. I believe that should 37 go forward and start being implemented, the very first question the board's going to ask is, what are the SLAs and who do they come to?

They're going to come back to this group, so there's not another group they're going to go to and say what should these numbers be. It's going to be this group. So we should define what those numbers are. Maybe we define them to say it's the same metric, or the same metric with different thresholds because they're tied differently or monitored differently, one [inaudible] Internet, one monitored locally.

I don't know, but we should address that, because it's going to be a question that's asked to us.

UNIDENTIFIED MALE: Sorry, just a quick note on the one set of numbers. I think there are two sets of numbers for when we get to thresholds.

BRAD VERD: [One set of metrics..]

UNIDENTIFIED MALE: Okay, good. So one set of metrics which might be split into two different thresholds. I just wanted to make sure we weren't thinking that there's just one, because I think there is a different one between – if there's going to be money handed out and whether you can even enter or get thrown out of the local root server system. I think that those will – money will be higher than if you're in.

LARS-JOHAN LIMAN: I agree. I suggest that this work party doesn't focus on the money [inaudible] bilateral problem to be solved by others.

FRED BAKER: I think that on metrics, [inaudible] similar to the way [inaudible] think about metrics. A metric is a way to measure something. Oh, and if the number should be something and you get some money as a result, [there's a policy.] I would really rather in this group talk about metrics, how we measure things, sort it out, and let thresholds [and] money be a separate question.

LARS-JOHAN LIMAN: I agree. We should first define the metrics as you know mentioned. I think that this group could do good work when it comes to defining thresholds for general service, and I would like to leave that contract part outside. But I actually [inaudible].

FRED BAKER: Certainly, we're going to need to do that and I have no disagreement on that point. I'm just noticing that at this point, we've gone 25 minutes talking nothing about metrics [inaudible].

DUANE WESSELS: Remote, you can still hear us; correct? I've lost Internet connectivity. Just checking that you're still there.

UNIDENTIFIED MALE: Yes .We can hear you.

UNIDENTIFIED MALE: I can hear you.

UNIDENTIFIED MALE: We're here.

DUANE WESSELS: My Zoom application quit, so if there's a hand, I don't see it right now.

UNIDENTIFIED MALE: There are none.

DUANE WESSELS: Okay. Yeah, so this is frustrating. I think that one of the reasons that we're sort of stuck here is it goes back to this issue of – it comes up in the context of correctness that the location of the measurement matters, at least for something like correctness.

So we can define metrics. I guess we could take the work party in a direction where we define metrics and just ignore the question of location. But [essentially,] someone has to answer that question at some point, somebody has to decide, okay, you do this measurement from here or from here and look [what the consequence is.] You know what I mean?

So for example –

BRAD VERD: The first question was research or SLE numbers, and [inaudible] objection to the proposal [inaudible]. So it sounds like that's been answered [inaudible].

FRED BAKER: [Well, Wes isn't here.]

BRAD VERD: So there's no objections, [inaudible] everybody in agreement [inaudible]. Second piece, the location piece. I guess I don't necessarily – I'm sorry, I missed some of the meetings. I've listened to a bunch of them. I don't follow the correctness [inaudible].

I understand that someone can't be responsible for the Internet [inaudible] monitoring probe, and then I hear all these discussions about gaming the system. I'm kind of like, [if they're going to game the system, good, that just means more root servers.] I don't understand why there's all the discussion. I guess talking about gaming the system. [inaudible]

FRED BAKER: [inaudible]. You have proposed that there be like ten probes.

BRAD VERD: Right, that was just me kind of throwing something out there.

FRED BAKER: If there are ten probes and all of a sudden those measurements are relative to those ten probes, everyone's going to want to move closer to those ten probes. That's what we're talking about with gaming the system.

RYAN STEPHENSON. Sort of. Because then you lose an aspect of diversity.

BRAD VERD: No, I disagree. I do not understand the diversity discussion around that whatsoever. So I would love to have somebody explain it to me. Maybe we can do it offline, but [inaudible].

UNIDENTIFIED MALE: Two things. One is on the gaming. I'll talk because Warren was the one who brought up gaming. I talked with Warren about this, and one thing that's really important to remember is that wherever we put the probes and however many we have, they are always going to – and probe number 12 is now looking at [D root]. It is going to go only to the closest instance that some routing mechanism has found for it.

So with the idea of a smaller set of probes, you don't have to put all of your instances near them. You just need to have some that are going to be near – and maybe just one that is going to be near enough to the collection of probes to get above the line.

So in the sense of gaming, what Warren was saying was he wants the only result published to be yes/no, not closeness, because if it's yes/no, people will be able to put a probe or two – so even if there's an operator whose purpose is to have a zillion probes in obscure locations in order to help people get root servers there, as long as they have one or two probes, one or to instances near where some of the probes are, they're probably going to look okay. So that's one level of gaming.

But the other one we need to remember is quite frankly, RSSAC just put out RSSAC 42 that says every operator gets to decide how they want to do things. Now, I think with respect to metrics, that means you may decide to never put anything that's going to get you good metrics.

But we are sort of saying if you still want to be a root server operator, one of the thresholds is going to be you no longer can be [in service,] which sort of [scoots on] RSSAC 42 in the sense that we're saying you can set up your network however you want, you can put the probes wherever you want, as long as you also have enough to be able to pass the minimum threshold for being a root server operator.

And that might be okay.

UNIDENTIFIED MALE: [inaudible] can put probes. I don't understand that.

UNIDENTIFIED MALE: I'm sorry. [I'm mistaken.] Well, we also have a you can put probes question with close versus far probes, but that's different. You can put instances wherever you want as long as you –

UNIDENTIFIED MALE: [inaudible].

UNIDENTIFIED MALE: And even with what I just said, [is this a bad idea?]

UNIDENTIFIED MALE: [inaudible].

UNIDENTIFIED MALE: [inaudible].

UNIDENTIFIED MALE: [inaudible] moved around to make it work?

DUANE WESSELS: I'm going to try to get us back on our little agenda. So if we can take that first question as resolved about purpose of metrics, then the next question in my mind is something we talked about briefly here, which is who operates the probes and sort of should they be operated by an independent third party, or could they be self-operated, self-reported by the operators?

UNIDENTIFIED MALE: [inaudible] metrics, whatever platform you use, [inaudible] different. But whether that should have [inaudible] taking the measurements or you have to [inaudible]. It's just a different way to measure, and you measure a different thing. [inaudible] know.

DUANE WESSELS: Right. My question is – let me [put it] this way. Should the work party be concerned about that difference between Verisign measuring its own root servers or [inaudible] measuring its root servers? Are those equivalent? Do you trust them the same way? Do they have the same meaning, the same weight?

UNIDENTIFIED MALE: So part of what Daniel's asking is actually quite relevant to if some of the measurements are going to have near probes – we've talked about the correctness one. By the way, I think if we're going to have near probes for correctness, we also have to have them for [inaudible] because if someone can intercept, put in bad data for correctness, they can and put in bad data for freshness.

So if we're going to somehow have near probes, those sort of inherently have to be run by the root server operator themselves, because they're going to know what's near.

So we either have [inaudible]. Duane, can you –

DUANE WESSELS: I did.

UNIDENTIFIED MALE: Now you have. Or now you fell off. Either we're going to have to have two sets of operators – the root server operators themselves doing the thing that we call near, and someone else doing all [of them.] Or we go to just one set of measurements and we don't worry about that near, far. And in the cases where somebody has failed a threshold due to correctness or staleness, we then have a second layer of let's figure out what happened and stuff like that.

In my mind, having one set of metrics that are all done far would be good, because also, we the near one, we have the question of L root, we have over 100 instances. Do we have to put 100 close ones, or do we

just do one, at which point we could have just done one with a far one because it's pretty close to one of the probes anyways?

DUANE WESSELS: Yeah, I see that is related for sure.

UNIDENTIFIED MALE: So I would hope that we would go towards one measuring organization doing the collecting. And we could plot it or whatever, but – and I would hope [that they'd be open on how they're doing it.]

DANIEL MIGAULT: [inaudible] I think we have to [weight the different results.] One is going to be more accurate [inaudible]

UNIDENTIFIED MALE: Exception.

DANIEL MIGAULT: No, [inaudible] own system.

UNIDENTIFIED MALE: Oh, self-assessment.

DANIEL MIGAULT: Self-assessment is probably the most accurate thing. The problem we have is [inaudible]. So the far away measurement is [inaudible] and I think we can [do it with the far one]. Given that, okay, it's just an indication of maybe something [is happening. It doesn't mean it has happened.] [inaudible].

DUANE WESSELS: Go ahead.

UNIDENTIFIED MALE: Sorry. I just think that for transparency it seems like it should be third-party so we don't get any kind of accusations of the fox guarding the henhouse.

DUANE WESSELS: Are there any [inaudible] – are there any that would be okay with being tasked with – in addition to your root service, you now also have to operate this probe infrastructure, you have to be responsible for maintaining these devices or these VMs, these applications, whatever? Is that a burden you're willing to [inaudible]? [Brad?]

BRAD VERD: [Are you accepting money?]

DUANE WESSELS: [Yeah.]

BRAD VERD: [inaudible]. That's why I just said SLA numbers and general service numbers. If you're accepting funds, then yes, you accept the burden. That should not be a question.

DUANE WESSELS: Any different opinions? Jeff?

JEFF OSBORNE: [Just at a high level] [inaudible] this organization [inaudible] converge on things. And the problem is money is a single [inaudible] transaction that changes everything. You can't converge on whether this is funded or not. If it's funded, it goes this way. If it's not, it goes that way. It can't converge on those two things and then hope we get money later. It changes everything.

LARS-JOHAN LIMAN: I agree with Brad [inaudible] obviously, [inaudible]. I think [inaudible] these boxes and adapt that to whatever these metrics [inaudible]. It's probably not going to be a huge effort. So in order to create some transparency and participate in the big picture, that's a reasonable request, if the amount of [inaudible]. There's a lot of [balancing] there, but I see that as in general, [inaudible].

FRED BAKER: So Brad, I agree that if money is in the discussion, then the guy accepting money is going to have to do something. My question is implementation. Does that mean that I have to double the amount of hardware that I have out there? Can I put it in a VM [inaudible] in the same chassis, run it on localhost, something like that?

I'd like to believe that I can implement that in any number of ways that may not have a [cost and operational burden.]

UNIDENTIFIED MALE: [inaudible] with what Fred said. I think running – being required to run probes makes perfect sense so long as the probes can be run in various infrastructures and not require hardware dedicated probes that are located in certain positions in the network.

KEN RENARD: Same thing. If we can build our own measurement, it's fine.

LARS-JOHAN LIMAN: [inaudible] deploying new hardware, [inaudible] very problematic unless we get paid for it.

BRAD VERD: I think we're [overengineering this,] cart before the horse. [inaudible] no offense, I understand there's certainly going to be an issue, but if there's a cost associated [and we're] getting money, then that should be

... tied to it. Regarding the independence of operators, if they want to do it anyways, they can go do it anyways, but they're not getting funds.

But I feel like VM, boxes, whatever it is, we can talk about that [inaudible] metrics [inaudible].

FRED BAKER: Well, if you're agreeing that's [inaudible], then I think you've answered my question. I heard your comment was yes, there absolutely will be a separate measurement infrastructure.

UNIDENTIFIED MALE: I didn't hear that.

UNIDENTIFIED MALE: Okay.

BRAD VERD: I'm not sure what you're getting to with that.

FRED BAKER: [A separate piece of hardware.]

BRAD VERD: Doesn't have to be. Could be a VM.

FRED BAKER: If you're happy with it being a VM –

BRAD VERD: It depends on your network. Your network could be one box that does routing and everything, or it could be thousands of boxes that [maybe a VM doesn't sit outside the border.] It could be anything. I can't tell you what that is. I can tell you what I feel the metrics should be and possibly what the threshold should be, and if we're exchanging money for an SLA, then I'm going to say, what are you spending the money on? And [inaudible] SLAs [inaudible]. If I don't agree with it, then we might have a conversation and we work out an agreement on what would be acceptable to me. That's the bilateral piece.

DUANE WESSELS: Are there any hands on Zoom?

UNIDENTIFIED MALE: [inaudible] I'm going to suggest that RSO are able to publish their own measurement and another party can confirm or infirm or [send new data]. If an RSO is not [inaudible] provide the measurements in a very simple way, I think we have the wrong measurement [inaudible]. Because I'm sure when you're monitoring [inaudible] adding the raw metrics [inaudible]. So that should be balanced with [inaudible] measurement.

DUANE WESSELS: So we have some metrics today. We have RSSAC 002 which is metrics in a sense, which are all self-reported. So this already sort of happens.

UNIDENTIFIED MALE: [inaudible].

DUANE WESSELS: But we don't publish availability and we don't do latency. And those are very hard to self-report, I think. In my mind, those are things that really make sense to do from an independent third party.

So the thing I sort of struggle with here is we've got two types of metrics in this work party. We're talking about [inaudible] really make sense to do from a third party, and [inaudible] really make sense to do self-reporting? And maybe that's okay, but it complicates our work a little bit.

Maybe those self-reported ones we should maybe just move into RSSAC 002 space and not be talking about them in the same context as the other ones.

UNIDENTIFIED MALE: I hate to poke at [inaudible] but BPQ, we've sort of just decided is going to be self-reported. We've been talking about the four major metrics, but then we also say "And maybe BPQ." We might want to throw BPQ out of this even though it is part of 037, because again, for simplicity, the four that we have, people seem to have agreed on.

And quite frankly, I think that for correctness and staleness, measure them from the outside. I don't see a reason to make additional self-reporting at this point for these metrics. I don't see a purpose to that. If there are things that affect [the ability to get money] or whether you're a root server operator, whoever's doing the measurement outside, however they're doing it is brought into question, then we will get those. And those would only really, I think, be - [inaudible] would be correctness and staleness. I don't think you could question whether somebody is measuring latency correctly from the outside if they've been doing it the same for the last [ten years] or whatever.

So I would say just let somebody outside do it and don't even bring any questions. I know there are some root server operators that say nothing can be added inside of our network. So if we're going to have them put something just outside their network anyway, it's okay if it's another 100 feet outside their network. And therefore they don't have to do anything.

DUANE WESSELS:

I agree that it would make our jobs a lot simpler if we could just focus on one of these, if we could say – my personal preference would be to say let's have this all be external third-party measured. [inaudible].

Hi Wes.

WES HARDAKER:

Hello. Sorry I'm late.

DUANE WESSELS: It's okay. You've [probably heard all the discussion] [inaudible].

UNIDENTIFIED MALE: [inaudible] a couple of times.

DUANE WESSELS: So we're talking about – a couple things [inaudible] together. Who should be responsible for performing measurements? Self-reported versus third-party? And a little bit – this gets into locations.

And we've kind of gone around and I guess I'll state my preference again. I would prefer that if we could focus on measurements done externally by a third party and [inaudible] even for these ones that ideally would be better measured by near probes for now ,we can not do that. [inaudible] the work party, and maybe we try and see where it leaves us, see if there really are issues with the correctness measurements.

UNIDENTIFIED MALE: I totally agree. I struggle in my head sometimes thinking that maybe just if we use Anycast to our advantage and put resolvers [inaudible] we put [inaudible] deal with that provider or move that, or decommission [that node.] So as long as [those rituals] aren't set to zero, long as they allow for some brokenness occasionally somewhere, [inaudible] third parties with correctness [inaudible].

DANIEL MIGAULT: The problem is that the measurement [inaudible] something can only say [it works well.] If it doesn't work, cannot say it's because [inaudible].

I believe it could [inaudible] measurements [inaudible].

DUANE WESSELS: You think it's different for a third party versus –

DANIEL MIGAULT: [inaudible].

DUANE WESSELS: [inaudible].

DANIEL MIGAULT: For example, [inaudible] I don't have the right data, and I have no way to say it's because you're serving the wrong data.

DUANE WESSELS: Can you speak up for the [inaudible] and also identify yourself?

DANIEL MIGAULT: Okay. This is Daniel, and what I'm saying is that some of the measurements I believe are going to be hard to provide some kind of conclusions about what is being measured. We should not have [inaudible] measuring from [third-party] correctness. I don't have the

right data. They expect the data. [How can I interpret that?] [inaudible].
I think that is a tradeoff [either way.]

DUANE WESSELS: [Jeff.]

JEFF OSBORNE: I wonder [inaudible] addressing this. You don't need to [inaudible] go to the three companies and see which one [inaudible]. If you're talking about third parties to measure this, there are an infinite number of third parties capable of doing global measuring of this stuff. Off the top of my head, [inaudible] Atlas probes, ThousandEyes, there's ample organizations capable of this.

And it almost seems like we ought to say to them, "Hey, what kind of useful metrics can you guys measure?" Rather than [inaudible]. It's just a thought. If we are doing outsourcing.

LARS-JOHAN LIMAN: I actually strongly disagree with that. [inaudible] can be measured rather than [inaudible] decide what's useful for us after that's measured is a totally different approach and I really don't like the one that you proposed. [inaudible]. We should decide which metrics we want to be measured [inaudible], because we want to be on top of the entire system.

UNIDENTIFIED MALE: It just strikes me [inaudible] another six months to simply see this offered by a commercial [inaudible].

LARS-JOHAN LIMAN: To use that as a start point, fine. To use that as the defining thing, not so happy. [inaudible] agree with Daniel that there's a big difference between having a [inaudible] close to the server. But I also think that we might not have to dwell into that because if you only measure from [inaudible] probes and you see an issue somewhere, then you go into that and investigate, and then you can use all kinds of tools to try to figure out what's going on and where the problem is. That should be [inaudible] so that people can [inaudible] themselves what's going on and whether this is the fault of the server operator or intermediate network or the probe.

So we may not have to dive too deep into [inaudible] measures, because that's part of the [inaudible] process later on. We might get away with just focusing on remote measurement.

PAUL HOFFMAN: I totally agree with Liman, especially – and very much disagree with Jeff about us going out and asking. We already have spent a lot of time sort of coalescing on when should measurements be done, how often and such. And these are being reported once a day. And as long as the SLA/SLE is not saying “If you blow it one day, you don't get any money or you get thrown out of the root server system,” that gives every operator an opportunity to go, “Oh, I failed for some reason yesterday and I want to now look at that” and such like that.

I think that as long as there is a fair [inaudible] recovery time between individual not passing a threshold and being able to [fix it,] I think that [this is a system sort of like what has already been ascribed] [inaudible].

DUANE WESSELS: So it's 9:00. This meeting is scheduled to go until 9:30, which I think is also the time that the [inaudible] sessions are [inaudible].

UNIDENTIFIED MALE: [They start at 10:00. We gave ourselves time.]

DUANE WESSELS: Okay.

UNIDENTIFIED MALE: [inaudible].

DUANE WESSELS: [inaudible]. So [inaudible] getting ready to leave. I thought maybe some others were going to –

UNIDENTIFIED MALE: [inaudible] 9.00.

DUANE WESSELS: If people need to leave –

UNIDENTIFIED MALE: [inaudible].

UNIDENTIFIED MALE: [inaudible]

DUANE WESSELS: So I'm still not on the Zoom room. Russ or any questions from the room?

UNIDENTIFIED MALE: No hands.

DUANE WESSELS: Okay. So I feel like we made a little bit of progress here on some of these issues. I can summarize and just make sure that we're all on the same page still. We've agreed to sort of focus on the service-level aspects of the measurements, and it sounds like we have a way forward, agreement that measurements will be done from third-party locations, not necessarily by root server operators or near probes.

And given that, I guess one of the next questions that we should stress is the coverage aspects. How much of the system do we want to cover with measurements? Or should the work party say anything about this?

UNIDENTIFIED MALE: [inaudible] what system?

DUANE WESSELS: The root server system. So the root server system is about 1000 different instances, and Paul and I have been doing some measurements and analyzing data and looking at this a little bit. [inaudible] measure from N point, how much of the root server system do you - [is this] in scope for the work party, or should we just only be defining how [inaudible]?

LARS-JOHAN LIMAN: If we go for the path [inaudible] of the actual number of instances. That's unavoidable. So in that case, it's just a matter of deciding how much of this do we need to cover. [inaudible]. I think that's the right way to go.

And I could probably [inaudible].

DUANE WESSELS: Russ?

RUSS MUNDY: I just want [inaudible]. When you say service level [inaudible], are you talking about as the RSS as a whole or individual RSOs?

DUANE WESSELS: Both, kind of. But I think we have to acknowledge [stats.] So this work party, as much as we can, we're going to avoid SLAs, [inaudible] talk about SLAs [inaudible] talk about service level expectations. When it

comes to SLAs, [inaudible] individual organizations can be held to SLAs.
You can't hold the entire root server system to an SLA.

But you could still, I think, state your expectations of the root server system and its [inaudible].

UNIDENTIFIED MALE: [inaudible].

JEFF OSBORNE: Is it really [inaudible] starting point of each of those [12] discussions?

DUANE WESSELS: Yes. Absolutely.

FRED BAKER: If that's not true, we've got a problem.

DUANE WESSELS: [That answered your question, right?]

UNIDENTIFIED MALE: Yeah, I just wanted to know what level we were going.

DUANE WESSELS: Okay.

UNIDENTIFIED MALE: [inaudible] things like correctness, [inaudible]. The other possibility is requiring some root server operators [inaudible].

DUANE WESSELS: Just [should probably say] more about what I was thinking. [inaudible] the work party should say something like any system that gets built should probably try to measure 25% or 50% or whatever, but I think 100% [inaudible].

UNIDENTIFIED MALE: [inaudible] you need 120 different physical locations [inaudible].

DUANE WESSELS: [At least.] Yeah.

LARS-JOHAN LIMAN: I disagree with your idea that what we provide [on NetNod's] side is the [cloud] service. We don't provide 55 different sites. We provide [a cloud] service. So that's what you should measure.

UNIDENTIFIED MALE: Also, we've got a couple of root server operators who are using commercial cloud services now. [inaudible] how the bleep can you expect any measurement correctly on those? And currently, there's only

one cloud service provider doing it, but other RSOs might pick other competitors [inaudible].

And I know I came in sort of late on this, but I've been going by what you started with the work party, which is per root server operator, never looking at instances. That's, to me, easy to look at. Anything beyond that where – for example, let's say that you said, oh, 10%, so each root server operator who has more than ten instances says, "Okay, here's [inaudible]."

They can't move them easily. That I think is a bad idea.

DUANE WESSELS: Yeah, that's not exactly the direction –

UNIDENTIFIED MALE: I know. But I think that's exactly where – that's not even all the way down the slippery slope. I think halfway down you're going to pass the sign of "And you need to not move the ones that you told us we should be measuring here." And I think that's dangerous.

DUANE WESSELS: [inaudible].

UNIDENTIFIED MALE: [inaudible] the conversation always goes back to the [inaudible] metrics [inaudible] fall into this SLA [inaudible]. I agree with Liman that the cloud service [inaudible] every instance I [have, that's unreasonable.] I

think it's reasonable if I'm accepting money to move forward [with] every instance I have.

So this is the difference between if you accept money, these are the obligations you sign up for. So I feel like what I hear the work party saying is we don't want to touch that. We want to touch money, we don't want to touch the SLAs or the SLAs discussion, we want to [inaudible] SLEs. But I feel like you're having that discussion. You're having an SLE discussion, which is what the paper's going to be, and I hope in the paper is, "In our discussion, we have identified a number of [inaudible] what could be or should be SLAs and here's what our thoughts were."

I just feel if they're getting money, they should report [inaudible].

DANIEL MIGAULT: Do we to define –

DUANE WESSELS: Daniel, speak up [and say who you are.]

DANIEL MIGAULT: What I'm [inaudible] about the SLA thing is that what we're defining in this work party is this is what we expect, and [inaudible]. If you go over that bar, it means something goes bad somewhere.

SLA is up to what the person is willing to do. It's an agreement and negotiation between what the person is willing to give you and expecting from you. And [inaudible].

UNIDENTIFIED MALE: [inaudible] coming back to your comment, how many TLDs were added in the last round? Let's say 1500. There are [not] 1500 different SLAs with each of those TLDs. They're not all different. There's one set of SLAs and one set of metrics for all 1500.

So I disagree. But –

LARS-JOHAN LIMAN: I continue from there. That type of SLA is [what I expect] people [inaudible]. That would be the common denominator for all the root server operators. [inaudible] someone is willing to accept money to do something better, that is specific negotiation between two parties that are not part of [inaudible].

BRAD VERD: Sure. The key point to what you just said is “do something better.”

LARS-JOHAN LIMAN: Yes.

BRAD VERD: If you accept funding to do what's expected; what is expected? This is what we're defining. We should define what is expected. And then if you want to do something better, great. But what's expected?

LARS-JOHAN LIMAN: To me, it's not expected to report [metrics] from every node [in the baseline requirements. If I get money, maybe it is, because if I get money, that's a negotiation. And the negotiation could lead to [inaudible] according to standards. [inaudible]. It could be each report [inaudible] normal standard and report [inaudible] different sum of money. It could be "Please perform to a better standard [inaudible] and report [inaudible] for this amount of money." So that's negotiation for me. The baseline should net require [inaudible].

if you want to do something better, [inaudible], yes, definitely money involved, or turning around – if someone wants to pay me for not reporting, fine.

DUANE WESSELS: Wes.

WES HARDAKER: [inaudible]. If we say we want [inaudible] cloud, you can't talk about 10%, because 10% requires 10% of the [number.] So you have to actually immediately [inaudible] infrastructure of how do we measure [inaudible] the world, and each proposal would likely have a different mechanism for saying this is how [inaudible].

The second thing is [inaudible]. Anyway, about the purpose of this group, right, this is the shepherd speaking. Our purpose is to get ahead of RSSAC 37. And Brad pointed out there's one SLA for all the TLDs. It's exactly accurate. And if we look at the model of what we expect might come out of 37 and [inaudible], the TLD operator specification says 500 millisecond [minimum round trip averages, you could do that from anywhere on the planet.]

The [inaudible] if we expect the service to operate as a whole, this is the minimum, and then if there is – individual operators have better ones. I expect that there's other SLAs between the TLD [inaudible] and the registry that they outsource their service to.

So there's probably actually more than just [inaudible] for RSOs. [inaudible] ICANN, because if we talk about the past, during all of the actually 37 discussions [inaudible] diversity. We expect diversity [in the same] SLA. [inaudible].

LARS-JOHAN LIMAN:

I agree with you about the percentage and the numbers. You're quite correct. So I think we need to do that from the other [inaudible] how many vantage points do we need to build a picture that is good enough for reporting it as a measurement? Do we need 15 [vantage] points? Do we need 25? Do we need 200? That's the way to approach it, not how large fraction of the system [inaudible].

UNIDENTIFIED MALE: Is that up to this specification, or is that up to the person that's going to implement the specification?

DUANE WESSELS: I think it can be up to –

UNIDENTIFIED MALE: [inaudible] I'm curious.

DUANE WESSELS: It seems to me like it would be appropriate for the work party to specify some minimum number of vantage points.

UNIDENTIFIED MALE: [inaudible].

DUANE WESSELS: [inaudible]. Otherwise, if it's just open, then we have too much uncertainty.

UNIDENTIFIED MALE: And we know that one is not enough because it's a worldwide system. There are [inaudible] issues.

DUANE WESSELS: Anything on the phone? Russ?

RUSS MUNDY: Yes. So a little bit ago, we were trying to summarize points we'd come to agreement on, and I think we reached agreement on the fact that measurement points or vantage points will not be local, that they'll be remote, but we also talked about perhaps moving some of the metrics themselves out of the document we're working on and into the 002 document as some of the more self-reporting things that come out of the 002.

Did we reach agreement on making that move? And if so, which of the metrics would be involved in shifting over to 002 instead of the document we're currently working on? Thanks.

DUANE WESSELS: Russ, I think we did not agree to do that. I think that for now, we would keep all of the metrics that we're currently talking about in the document. The one exception may be the BPQ stuff. I'll sort of put forth the idea that maybe that should I guess be thrown out altogether.

UNIDENTIFIED MALE: Because [inaudible] BPQ would be derived from 002 anyway. Or if not 002, 002+, something like that. I think it's fine to leave that there, and maybe in our final report say, "By the way, there's another set of numbers that people have been collecting for a long time that may affect SLAs, SLEs," but we're not biting that off here because they're already there. Does that make sense, Russ?

RUSS MUNDY: I think so. I would maybe describe this as saying what we provide in our document we're working on is a description of how one goes about computing from the 002 statistics some kind of BPQ value, and that's all we need to say in our document. Is that accurate, Paul?

PAUL HOFFMAN: Yeah, that sounds fine. I think that we have enough already in 002 and long enough measurements where we could come up with something to say where that's valuable. Again, it's not the BPQ that's in 037, because 037's BPQ is also about capacity, but we can say here is some BPQ-ish things that are not about capacity that we already have.

RUSS MUNDY: That sounds good to me. Okay, thanks. Just wanted to get clarification on if something was moving or not. Okay, good. Thank you.

DUANE WESSELS: Mauricio, go ahead.

MAURICIO VERGARA ERECHE: Can I come back to [inaudible]? I think it's a good idea to cover [inaudible] but also need to have the investigation of they're not going to be everything all in the same place. [inaudible] something like that needs to be taken into consideration.

UNIDENTIFIED MALE: Yeah, I agree that's a very important stance. I would like to think that [inaudible] write that down. In the TLD case, they have something like that already. It says major geographic locations are [inaudible]. So we need to write something down like that.

UNIDENTIFIED MALE: So I'll take the point on at least getting not only what we said in the TLD [inaudible] – I know a bunch of people here know that because they're on the pointy need of it, but actually possibly technical description of how we, ICANN, are meeting that. I can do that within the next month, to get that. That may be good enough for [inaudible] but at least that's a worked out example that people in the DNS world know. [inaudible] find it out this week because Francisco's here this week.

MAURICIO VERGARA ERECHE: Also, maybe set out the [inaudible] how many instances are out there. Maybe that will be useful because that would give us a [inaudible] take into the future, how many [reports] do we need [to give information] [inaudible].

BRAD VERD: [inaudible]. I'm a little [discomforted] [inaudible] because like Liman said, it's a cloud service. It shouldn't be based upon the number of instances. [inaudible] number of instances where the optics are that "I need my own root server, I need [inaudible] million dollars" whereas the services available with 1000 [inaudible] what the message is. But the monitoring of it shouldn't be based upon [inaudible].

UNIDENTIFIED MALE: Let me just jump in, because [inaudible]. So we have 13 parts of RSO equivalentents now. At some point in the future ,we might have 26 because this has all been opened [and stuff.]

UNIDENTIFIED MALE: [inaudible].

UNIDENTIFIED MALE: Well, but let's go to 26. Do you think we need to double the number of probes if we have doubled the number of [operators?]

BRAD VERD: No.

UNIDENTIFIED MALE: Okay. That was a straight question of – and I see –

BRAD VERD: [inaudible].

UNIDENTIFIED MALE: Okay. What I'm saying is if we increase the number of RSOs, do you feel like we need to increase number of probes? And humorously, if we reduce the number of RSOs, should we throw out probes?

LARS-JOHAN LIMAN: Wrong angle. You need to focus on how big is the Internet.

BRAD VERD: Right. That's what I was saying. If you look at it from – that's why when my comment of ten that was said weeks or months ago was pick the ten largest [inaudible] exchanges or Internet exchanges in the world, [inaudible], and that's going to cover a vast majority of [inaudible] Internet. And that – I think what I said [in the] conversation was “And that number should be reevaluated every couple of years.” Because Internet exchanges change in size, new ones pop up, things move. That should be reevaluated. But we shouldn't have an equation that is one day we're going to have so many [probes] out there. That's [just scary.] And I don't think [inaudible].

Because one would say, how big is the Internet? We're talking about a number of different characteristics. It could be - [inaudible] really short, then presumably you'd need fewer probes all over the place. [inaudible]. Whereas [inaudible] have these special cases that we have to measure, I think [inaudible] different segments of the Internet [inaudible] latency, [inaudible].

LARS-JOHAN LIMAN: Sure. Absolutely. [Shifted the picture from the angle I was looking at] and that [inaudible] the first one to make you shift the entire thing around. So you're absolutely right.

DUANE WESSELS: Alright, so we've only got a few minutes left. I wanted to talk just briefly about or mention briefly that – I think [inaudible] mentioned this already, he's sort of done a proof of concept implementation of the metrics and he's running it from his DigitalOcean datacenter.

UNIDENTIFIED MALE: You said ten. DigitalOcean happens to be in 12 datacenters, so I'm just stuck in [inaudible]. I'm running it now – I'm actually not reporting every day, but I'm reporting exactly the way that it looks in the – [inaudible] current [inaudible]?

UNIDENTIFIED MALE: Yeah.

UNIDENTIFIED MALE: And so I can send it out to the list or whatever.

DUANE WESSELS: Yeah. [If you have it,] post it to the list.

UNIDENTIFIED MALE: I have. I can send it to the list. [inaudible] only have a few minutes, but basically, what I'm seeing that might be of interest is day to day, some days there are a bunch of connection errors, but the one I could show you today – because I actually redid the metrics the other day. I lost all my old stuff. There were no connection errors.

Generally, what we're seeing is that for an individual RSO – just eyeballing – some days your TCP v4 sucks relative to what you would want, and other days it's just fine. F and G basically have zero latency because of – I'm assuming because of cloud. Theirs is so much ridiculously lower than everyone else's it just looks weird.

UNIDENTIFIED MALE: In our private conversation [you had a] question to me that made sense, [separate out these metrics by] [inaudible] TCP. So [inaudible] take that to the list.

UNIDENTIFIED MALE: [inaudible]. Now, would it bug you if I sent that [inaudible] included? Right now, I collect the metrics to the one central place and then I run this [show me,] and then that [show me] has an additional section of [inaudible] that are [current.] So I'll put that out there as well. But if you collect some of the v4/v6 stuff, the numbers in my mind look a little bit better. So instead of one number for v4 TCP, one number for v4 UDP, one number for v6 UDP, one number for v4 TCP, stick them all together, because from a resolver's point of view, it's going to pick the best of those four anyways.

So I'll send that out. And again, this is not to say that Brad's view on 10 or 12 is right, but this is a way of at least looking at what happens when you do – and they're not all at ISPs, but they're at major datacenters, and a bunch of them are in other places. There's one in Bangalore, there's one in Singapore, there's a couple of Europe [inaudible]. So just

a way of looking at that smaller number than the Atlas probes. [So I'll send it to the list]. Thanks.

DUANE WESSELS: So let's wrap it up. Any last comments from the Zoom room? [inaudible] participants before we -

OZAN SAHIN: I see no hands. Put up your hand if you want to make a last comment. No hands are waving.

UNIDENTIFIED MALE: [inaudible] the resolvers are going to pick the last one. [inaudible].

UNIDENTIFIED MALE: [inaudible].

UNIDENTIFIED MALE: [inaudible].

UNIDENTIFIED MALE: They will pick the best one, but they will probably pick one for at least six hours.

OZAN SAHIN: Okay. Thank you everyone for making the call, and [inaudible].

DUANE WESSELS: Yes. Thanks, [everyone.] Thanks for the people on the phone who've [inaudible]. We'll send out a picture of how funny this room looks.

UNIDENTIFIED MALE: [inaudible].

UNIDENTIFIED MALE: [The handles on the chairs.]

DUANE WESSELS: Yeah. Okay, so actually, Andrew or Ozan, are you [inaudible] the meeting?

[END OF TRANSCRIPTION]