

---

KIM CARLSON: Welcome to today's RSS Metrics Work Party call on August 8th at 17:00 UTC. On today's call, we have Duane Wessels, Wes Hardaker, Jaap Akkerhuis, Karl Reuss, Ken Renard, Mauricio Vergara, Paul Hoffman, I presume, Russ Mundy, Shinto Sato, [Tale]—I'm not sure who that is—Robert Story, and then there is one other participant who is on the phone only. If you can make yourself be known now, and I can accurately add you to the attendance.

PAUL HOFFMAN: I am the phone person with the 831 area code.

KIM CARLSON: Okay, great. Thank you.

DAVID LAWRENCE: The [Tale] was David Lawrence.

KIM CARLSON: Okay, thank you. And from staff, we have Steve Sheng, Andrew McConachie, and myself, Kim Carlson, pinch hitting for Ozan today. To my knowledge, we have not received any apologies, but I will update the official attendance after confirming. As a reminder, these calls are recorded and transcribed. Please also remember to state your name when speaking for purpose of the transcript. When not speaking, please keep your phones and microphones muted to avoid background noise. And with that, I will turn the call over to Duane. For those late

---

*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.*

---

attenders, I will make sure to add you to the attendance as well. Thank you.

DUANE WESSELS:

Alright, thank you, Kim. I sent out a agenda sort of late yesterday. If anyone has any additions or suggestions for the agenda, please speak up now, or maybe type something in the chat. Otherwise, I think we should get into it. I'm sharing my screen with the Google Documents. You're also welcome to, of course, join the documents on your own. The link is in the chat.

One of the first things I wanted to do today was just to quickly recap some of the discussions we had at our previous meeting, which was an in-person meeting during the IETF in Montreal. The work party met for something like an hour, hour and a half. I feel like we made some important decisions, and I wanted to make sure that everyone was aware of those.

One of the decisions we made was that going forward, the work party and the documents should really focus on the service-level aspects of the metrics, and we would consider some of the other things that—we called them research uses, research aspects, or end-user behaviors—to be out of scope for this work party. So, that helps us focus our efforts here, and make some progress, and meet some of our deadlines. Any comments or questions about that decision?

---

WARREN KUMARI: I have a quick question. Apologies I wasn't there, but something I didn't fully get from the discussion was, are we planning on simply having a pass/fail metric, or are we planning on actually publishing the numbers themselves? I think they have very different implications between those two.

DUANE WESSELS: Yeah, we're going to talk about that later. That's still to be determined.

WARREN KUMARI: Thank you.

DUANE WESSELS: The other decision that we made, somewhat related to the previous one, was we talked about where probes might be located and who might be operating them. Previously we had put forth the idea that maybe, for some of the metrics, we would want the probing infrastructure to be very close to the servers, and that those would then out of necessity be operated by the RSOs.

But we took at as out of scope, so the decision was to not have any probing infrastructure operated by the root server operators, and that all the probes would be located in the network, farther away—at this point, not really saying where. That's still to be determined, but we did agree that they would not be operated next to the root servers or by the root server operators. Any questions or concerns about that decision? Okay.

---

Along those lines, I have been editing the document to reflect those decisions a little bit. Some text has been deleted. Some has been marked with a strikethrough, so to be deleted in the future. So, when you read through the document, keep those things in mind, I guess.

Also, some changes to the document that I have been doing are a little bit of restructuring and adding new sections. I don't think the table of contents reflects these changes yet, but there is now a Background and Scope section. This first paragraph is new, but the second half of this was moved from somewhere else. This was previously in the Related Work section. That's just a little bit of restructuring there.

There's this section 2.2, which describes the things that I was just talking about—uses that are not in scope for this work, or at least not in scope at this time. And then, this Relationship with Prior Work is the same. Again, it was in a different spot before, so that's been moved. The terminology is the same.

There's this new section about probes and vantage points. Previously, we had this part, which is now ... You can see it's stricken out. This was what we had before. So, this is just starting at restructuring that, and here or there, just placeholders for things that the work party needs to make recommendations on. Need to make recommendations on the number of vantage points, location of vantage points, and connectivity. We'll talk more about that at the end of the call, if we have time.

There was a suggestion one or two meetings ago from John Heidemann, which I think was a good suggestion, which was to move the example results to a separate section at the end. That has been done. You can

see here, where I'm scrolling in section four, there used to be a table of example results at the bottom of each of these, and now that's all the way at the very end, these tables. These tables still need to be tweaked. There's still a lot of things to be fixed or decided here, but at least you can see that these tables are all in the section at the end now.

Then lastly, the last thing I want to point out—changes to this document—is I've added an appendix, which is some work that I did using RIPE Atlas data. I don't know if this will necessarily survive to the final product, but at least I think that this will be useful in our discussion about numbers of probes, and locations of probes, and things like that. That is there for now.

That's what I wanted to cover, in terms of document restructuring. Any comments or questions at this time, related to that? Paul, go ahead.

PAUL HOFFMAN:

Hi. Sorry about the delay on muting. Related to the last bit, at the meeting two meetings ago, I had put up a table of some preliminary results I was getting from doing measurements at 12 locations and such. And I said I would send more to the list, which I haven't, but I will do that later this week, if that would help as well—sort of in the same vein of what you have there in the appendix, so that people ... I mean that it's easy for—it probably should be ripped out before the document is final—but for people to get a feeling for what kind of measurements are we seeing if we do limited—like using various kinds of limits. Yours uses a lot of RIPE Atlas probes. Mine uses VMs and data centers and such.

---

I'll send those results to the list, and then folks can decide if they want to put some ... I've been keeping them on a daily basis, because that's the way that we're suggesting in the document that they be kept, and that might be useful if people want it.

DUANE WESSELS: Yeah, thanks, Paul. I agree. I would say even add them to the document as another appendix. We should also probably add the work that [John Kristoff] has done for us. I think all that is useful in our discussions. We can decide later whether we want to keep it or remove it for the final version.

PAUL HOFFMAN: Okay, I'll put a recent day in, but I've been keeping them for many days. I'll send the text file of all the days, so people can maybe do some visual comparisons, but I'll pick a day and I'll stick it in the appendix later.

DUANE WESSELS: Okay, that'd be great. Thanks. Alright, I don't see any more hands at this point. If the hand thing is a problem, and if there's a silence—there's a gap—please feel free to just jump in. If there's no more questions about that, then I will move on to the next item, which is what Warren was just asking about.

I believe, Warren, your suggestion was originally raised in the context of our discussions about latency. We had been talking a lot about how publishing latency measurements or results could lead or encourage operators to game the system. Your suggestion was not to publish the

---

actual numbers, but to just publish a pass/fail, and that may alleviate the desire to game the system. This is something I would like the work party to discuss, and also to consider whether this is something we should do only for latency or for all of the metrics. Warren, did you want to say anymore, since I think this was your suggestion originally?

WARREN KUMARI:

Yeah, I was shouting at my phone, but apparently, I wasn't unmuted. Yeah, if I can just expand upon that a bit ... I think that in many cases, as soon as you attach a number to something, people automatically try and optimize that number along whether it's useful or not. There's a human tendency to do that. But also, other people looking at it, unless they've thought deeply about it, are going to assume that either bigger numbers are better or smaller numbers are better, but somehow there is something which can and should be optimized.

So, I think, unless there is a good reason, we should set a ... This is what we expect as a service level, and then we should either just say how often or what percentage of the time people are meeting or violating it, or something very similar. What I don't want to have is a set of arbitrary metrics, and then people start trying to optimize, just to make the metric look better.

DUANE WESSELS:

Okay. Yeah, I think that's a good suggestion. I think it's a path forward, and I think makes a lot of sense for the latency one. One question I had about this, though, was if you think about something like the correctness metric, which in our discussions, we've talked about holding

---

correctness to very high standards—maybe even 100% correctness, right? If an operator fails the correctness metric, is it sufficient to just say that it was a failure, or would there also be a need or a desire to have a little bit more there and say what the failure actually was, like it was a failure to validate this type of record or something like that?

WARREN KUMARI: Yeah, if I can respond ... I think that for something like correctness, yes. That should definitely be at a much higher level. But correctness isn't really something where I generally see there being a number associated with it. I don't think that you're 63% correct versus 65% correct. I think you're either correct or not. When you're not correct, I think there should definitely be some explanation of in what way you are not correct, but I don't really see that as being a number that people would try and optimize for, other than 100%, all the time.

DUANE WESSELS: But it's more than just a pass/fail indication, right?

WARREN KUMARI: Yes.

DUANE WESSELS: Okay. Paul, you have your hand up.



---

PAUL HOFFMAN:

Yeah. I like Warren's idea a lot—that is, of not showing numbers. And quite frankly, I don't think we will ... We will see when we start talking thresholds, but for correctness, if we have a threshold of 100%, then that's the same thing. If they failed, people are going to see the failure. I believe, in previous meetings, we have said correctness is a different issue, because it's so freaking important. Maybe we have a thing that says ... Let's say that we pick 99.99%, and somebody doesn't reach 100, but has one incorrect one, I think it's okay for us to say, "And for that one, we will show what was incorrect." I also think that this group might pick 100% for correctness.

But going back to something you just said, Duane, which I think wasn't what I remembered, you said that we were thinking of doing this because of latency. I remember people also talking about this with the availability measurement as well, because availability will be slightly less than 100% in some cases, and such like that. So, I think we should also have this—just show the pass/fail for latency, availability, and staleness. Maybe correctness is different, and maybe correctness for every ... Even if we don't have 100% as the threshold, for everything below it, we show the public what was the incorrect value.

DUANE WESSELS:

I'm in agreement with you there. The part where we show what the incorrect value is, does that apply only to the correctness failure, or does that apply to other failures as well?

---

PAUL HOFFMAN: I would say it's only the correctness failure, because the other ones are all relative measurements. That is, latency and availability are based just on how long and how good was the SOA. In fact, so is the staleness at this point. We haven't completely figured out how we're going to do the staleness measurements. But none of those really have DNS data values associated with them, whereas correctness does.

DUANE WESSELS: Alright, thank you. Warren?

WARREN KUMARI: Apologies, again, if this has been covered and I just missed it. When we speak about correctness, different people have different views on what counts as correct. For example, I'm just putting in something at the moment from Geoff Huston on here, Scoring the DNS Root Server System. Are we only talking about, is the answer that you got what was provided in the zone file from the IANA, or are we also talking about some of the stuff which Geoff covers, which is did you set the fragment [bit] when he thought you should?

DUANE WESSELS: No. So, what we've got in the document is a definition of correctness that, at this point, only relies on DNSSEC. You do a query. You do different types of queries, and they all have to validate. One thing that we've talked about with correctness, that is not fully-described yet, is that you may want to do correctness checks over things that aren't

---

signed, like some glue records. So, we don't yet have a method for doing that, but it's still on the table.

WARREN KUMARI: Okay, and apologies for missing or forgetting that part of the document. But yeah, as long as it's what's in the zone file, I personally think you can check anything that's in the zone file.

DUANE WESSELS: Thanks. In the chat ...

PAUL HOFFMAN: Duane, did you see that ... Okay, sorry.

DUANE WESSELS: Yeah.

PAUL HOFFMAN: He was just going to point you towards that.

DUANE WESSELS: In the chat, Mauricio raised a question, which is something that I had also been wondering about. This gets to the extent to which you provide, maybe, more details to the operator than you would provide in a published report. I think that's a reasonable—maybe even also necessary—thing to do. I think if somebody operating these metrics is saying, "Operator X, you failed this metric," you really need to provide

---

the actual values, and more background on the failure. I think it would be appropriate to say as such in this document. Does anyone have a different opinion? Mauricio, does that meet your expectation?

MAURICIO VERGARA: Yeah, thank you.

DUANE WESSELS: Okay. Russ, go ahead.

RUSS MUNDY: Yes. Thanks, Duane. Although I'm generally in agreement with the statement that you just made, I guess the concern I have to an extent is ... The hard part might be what we actually write in the document. I think the idea of whoever and however is actually doing what I'll the official metrics—whoever ends up doing that. It's probably good that extensive information, and maybe all information, is available to RSOs, but only the publicly-aggregated information is what we say will be totally available. Yeah, the idea, I think, is good, that RSOs need to be able to say, and perhaps somebody might, or should be able to see everything that was measured, but we need to say it carefully.

DUANE WESSELS: Okay. Would you phrase it as, "See the information for their own operations?"

---

RUSS MUNDY: Yes.

DUANE WESSELS: Not all the information for all the operators, I guess, right?

RUSS MUNDY: Yes.

DUANE WESSELS: Okay.

RUSS MUNDY: Thanks, yeah. That was the thinking.

DUANE WESSELS: Okay. Another question I have on this pass/fail topic is do we apply the same to our RSS measurements—so, to the entire system? Should the published results say that the RSS passed or failed its thresholds, without saying what the actual numbers were? Paul?

PAUL HOFFMAN: I'm still skeptical that we're going to actually come up with RSS thresholds.

DUANE WESSELS: Well, let's assume that we do.

PAUL HOFFMAN: Okay. Assuming that we do, I would leave it as pass/fail, and at any time that we failed ... And again, it's not clear who's going to ... In the RSSAC 37 context, it's not clear which group is going to be evaluating those thresholds, but it would be up to that group to say ... The public could see that it failed from that group. It would be up to the group to say in what way it failed. I think that that's fine. If there is a group that is measuring—official group that is officially measuring these things—they should say to the public what has gone wrong. Again, I'm just fairly skeptical that we're going to be able to come to agreement on those.

DUANE WESSELS: Yeah, understood. And if we don't, then I guess we don't have this problem. Anyone else have thoughts on pass/fail stuff? Fred, in the comment, I see you have the question, "How does one pass without a stated threshold?" I think you need thresholds, obviously. I think the point that Paul was making was that while the group may come to a point where we make recommendations on thresholds for individual RSOs, we might have a tough time with thresholds for the RSS. I'm not as skeptical as Paul on that, but I think that was his point.

Alright, I think let's move on to the next agenda topic, which is reporting metrics separately for UDP, TCP, v4 and v6. The way the document is written right now, for some of the metrics, or for many of the metrics, there is ... The idea is that you report these separately for these different transports—in particular for latency, and availability, and so on—both for RSOs and for the RSS.

---

One thing to think of is that's a lot of numbers to report. People may be confused by the fact that you can pass in one transport and fail in the other. Other people may not be so confused by that. In my mind, if you don't separate them, then it's not clear. When you're doing the individual measurements, it's not clear which transport should be chosen, so maybe you should always do all the transports. But this is certainly open for discussion. It would simplify things, I guess, in some sense, if the transports were combined somehow. I see, Paul, you have your hand up. Go ahead.

PAUL HOFFMAN:

We had talked, both at the last meeting and at the previous meeting—at the previous IETF. Some people brought up the idea that what we're really trying to do ... We can do the metrics for all four, but should the thresholds be on the combined or not? That is, it certainly doesn't hurt to do the metrics. But now that we have agreed that this is going to be pass/fail, what would it mean for Q root to pass on latency for three of those four, for example?

The other thing that people said at both meetings, which I think is relevant here, is that the reason why we're having thresholds is to indicate what is acceptable to a user population. In the case, the user population, basically, are the resolvers that are using the root servers. If a resolver is talking to a root server, and one of those four doesn't work for it, it will switch. That's sort of the nature, and that's what we assume.

---

So, I think combining the four not only makes it easier for us to come up with a single metric. Also, if somebody fails that, then we can say publicly what it was. But it also reflects the way that the resolvers talk to root servers, I think—that is that they’ll pick v4 versus v6 based on Happy Eyeballs, where there’s not actual eyeballs here or whatever. They’ll pick UDP/TCP based on, “I like to fall back from UDP to TCP,” or, “Boy, I really like TCP.” I think we should assume that we’re trying to do measurements around what resolvers do, and I don’t think that they separate them out that carefully. Hopefully that’s helpful.

DUANE WESSELS:

Alright, there’s a lot of hands up, but I want to just respond to one thing, Paul. I think that when you’re talking about understanding what resolvers do, that’s getting into the part of the stuff that we’d said was out of scope for this work party. That’s the end-user experience, and that’s not what these metrics are about in my mind.

PAUL HOFFMAN:

Maybe I said it wrong. It’s what we would expect them to do, based on RFC 1035. Maybe not. Let’s see what other people ...

DUANE WESSELS:

Yeah, I think other Paul had his hand up first.

PAUL EBERSMAN:

Yeah, I think that from the operational sense, while yes, the protocol will work around the failure, knowing that v6 TCP fails or works versus



---

v4 is rather critical to the operators who need to fix it. I can't see having it conflated. We might have a separate thing that says, "Is there an answer that was pass/fail," i.e. did one of those four work, or did Happy Eyeballs or something similar work for availability? But I think we'd still need to have the separate pass or fails for each of the combinations of protocol and transport.

DUANE WESSELS:

Okay, thanks. In the chat, I think Ken has a comment that's along the same lines. Russ, you want to go ahead?

RUSS MUNDY:

Yeah, thanks. What I was thinking here, as part of this discussion, is at least for the different transports, it [might] be useful to have a single-end way to make it visible to the general population of the internet, what the state was. But instead of pass/fail, we might want to have a green/yellow/red kind of scale, where green was everything succeeded, red was one or two failed hard, and yellow was anything in between.

Of course, the information, like we talked about earlier ... If all the information were made available to the RSOs in question, then they could have the information. But for the general population consumption, if you will, the traffic signal—red/yellow/green kind of thing—might be better for some of these than pass/fail, especially for the transports.

---

DUANE WESSELS: That's an interesting idea. It would require some thought to how you would do that for each of the different metrics, I guess, because, of course, some of the metrics are not transports-based and some are. It would mean, of course, we need to come to agreement on what each ... I guess we would have to define thresholds for each color, essentially, or something like that?

RUSS MUNDY: Something like that, yeah.

DUANE WESSELS: Yeah. At first blush, to me, it seems a little bit more complexity than we wanted at this point, but I guess if someone's willing to write down some rules or whatever, then I think we could consider it. Warren, are you next?

WARREN KUMARI: I think I was. So, this entire thing, I think, is going to be one of the largest byte shedding events I can think of—choosing the numbers. I really like the idea of having a red/yellow/green type thing, but I think that the byte shedding we're going to have to just get pass/fail is going to be hard enough. Having the byte shedding for pass/fail/really badly failed is going to be even more fun.

DUANE WESSELS: You're up for it, though, right?

WARREN KUMARI: Eh ...

DUANE WESSELS: Alright. Paul?

PAUL HOFFMAN: Can you hear me or not? Actually, I'm just trying something different.

DUANE WESSELS: Yes.

PAUL HOFFMAN: Okay, great. That's cool. It doesn't say it in the phone. So, two things on this ... One is that ... Once I send out my test file of the results I've seen, we're going to see something that I think is going to bother some people, which is failures in v4 are actually higher than failures in v6 because you get some resets, and you get some routing errors. So, going back to what Warren just said, yeah, byte shedding that, and trying to explain why v4 might have a lower threshold might be interesting.

But I want to go back to what Russ was saying about red/green/blue as a thing that's going to the public, versus as what's going to the root server operator, who might not be passing. I liked our decision, less than half an hour ago, of pass/fail. I don't think that red/green/blue is actually going to be useful to anybody because it's not actionable.

---

If somebody who ... Even a resolver operator who happens to be following these, and sees that P root has oranges more often than others, what are they going to do about it? Are they going to somehow manually take P root out of their search path? No. They're just going to let the resolver do their thing. So, I am not as interested in going down red/green/blue by transport.

If we want to do different thresholds, that's fine, and then a root server operator might get thrown out of the root server system if their v6 transport over UDP isn't as good as it should be or whatever. I just think that that's actually not in the spirit of this. Letting operators know how they're doing? Absolutely. But that's different than doing this publicly, or possibly throwing them out of the root server system because they were orange too often.

DUANE WESSELS:

Alright, thank you Paul. Warren, did you want to talk about this giant link that you just pasted in the chat?

WARREN KUMARI:

Yes, and I hope this isn't off topic. I don't think it is. But I wanted people who haven't been involved in this for a while to understand some of the complexity around choosing the metrics and the numbers. Paul said that he sees more failures in v4. This is RIPE's DNSMON, which sees way more failures in v6. And I'm in no way doubting Paul's numbers, but I'm also in no way doubting RIPE's numbers. The trickiness here is exactly how one's defining what counts as failure, what doesn't, etc.

This is more just a cautionary note for people that doing network measurement metrics is really hard, and making sure that we choose the right metrics, and that they're perfectly defined, and that we have the methodology fully explained is actually really important. I think I'll get off my soapbox now.

DUANE WESSELS:

Thanks, Warren. Another thing interesting about the DNSMON page is ... They don't use pass/fail, but they use red/green, so it's binary. It's not ternary or something more complex. Russ, go ahead.

RUSS MUNDY:

Yeah, thanks Duane. Just a quick response to Paul. I'm not sure if I was unclear in what I said earlier, in terms of the red/yellow/green traffic light kind of thing. What I was thinking that that might be useful would be to combine—a way of combining the separate transports into a single representation that many people in the world are used to seeing, red/yellow/green, whether it's project status or coming up to a traffic signal at an intersection. It's certainly not anything I'm hard over about, but it's a fairly common representation, and a way to combine stuff.

DUANE WESSELS:

Yeah, alright. Thanks, Russ. I think we have good agreement that at least pass/fail is a very good approach. It sounds like we need to think some more about adding colors, or making it a little more complex. Maybe we'll see if somebody wants to come up with a specific plan around

---

---

that. Russ, is your hand up again, or is it leftover? I was going to move on to the next topic.

RUSS MUNDY: Leftover, sorry.

DUANE WESSELS: Okay, alright. Sounds good. The next topic is, as I showed in the Google Doc. I forget what section it was now. It's probably an unnumbered section. There's a section that talks about probes and vantage points. I feel like this work party should probably make some recommendations about numbers of probes, locations of probes, and other things like connectivity requirements. Strictly speaking, this may not be necessary, but based on the data that we've looked at—Paul's data, and even what Warren was just showing—we can see that the results that you get depend a lot on the number of vantage points and their locations.

I think it would be very good if we could make recommendations along these lines. I don't think we really want to do that today, but I would like to have some discussion on how to proceed. Maybe some folks would be willing to volunteer to take a stab at writing some text here. We did talk a little bit in Montreal about number of vantage points, I believe. And there were some pretty strong sentiments that the number should not be tied to the number of root server instances. So, any recommendation coming out of here should be a fixed number that doesn't have to constantly adjust as the root server system grows or shrinks. Personally, I could go either way on that, but that's a reasonable approach to take.

So, to open this up, does anyone feel like they have thoughts, or a good way forward on how the work party can make recommendations on the number of vantage points? Also, as I said earlier, there's some data in the appendix which can probably be useful here. I can talk about that, if we have nothing else to talk about. But Warren, I see your hand is up.

WARREN KUMARI:

Yeah, so I think that some of the number of vantage points depends, again, on are we only doing pass/fail, is the latency metric set, and where the latency metric is set. If the latency metric is set to one second, then having a very small number of probes seems reasonable. If the latency pass/fail is set to 5 milliseconds, I think you're going to need way more. I don't know who could really answer this exactly.

I did have a bit of a rant or argument with Paul Hoffman when I ran into him in Montreal—or maybe it was by phone—saying that I think we just need a very small number. And then, I discovered that, at least in his testing thing, the number of probes, or the cost to run a probe was really small, and the cost to administer the probes, if they're done through something like Ansible is also really small. That kind of changed my view from, "I don't see why we need more than four," to "Eh, if they're this easy to do, let's do a bunch more." And that was just a bit of background.

DUANE WESSELS:

Okay. Thanks, Warren. Paul, go ahead.

---

PAUL HOFFMAN:

I don't have a strong feeling. You said there were strong feelings. Brad's not on the call. He was one of the people with strong feelings. We should clearly do this on this list. I was reminded by someone else on ICANN staff that once we come up with a recommendation here, it's actually a recommendation to the ICANN Board. One thing that the ICANN Board understands is the ICANN regions. We've got five official regions. That's all discussed in the ICANN bylaws and such like that.

I'll send out to the mailing list a set of the maps of the regions. Since we know that number—there are five of them—I could believe that if the work party felt like 10 was an okay number of probes, and easy thing to do is to say two in each region. Or if the work party thinks 50 is the okay number, we could say 10 in each region. So, without having to go into more specificity of fast/slow and things like that ... And the regions are pretty big. That's just the thought. If people want to have some sort of geographic base, and aren't sure on how to do it, ICANN's already done it for us.

Like I said, I'll send out the maps. They are pretty sane, other than the fact that Greenland and Iceland are part of Europe. I think that's a political choice, not a geographic choice. Other than that, when you look at the map, you'll see, yeah, this looks divided reasonably-well, I think.

DUANE WESSELS:

Okay, thanks. Paul Ebersman.



---

PAUL EBERSMAN:

Yeah. Having gone through this exercise with ThousandEyes, using RIPE Atlas probes and a few others, and trying to do competitive performance numbers between various large recursive services, I learned painfully that network topology and geographic topology don't always map cleanly, though there are certain things that you can make assumptions about. Getting in and out of places like Africa or Central and South America are going to be a lot slower if you don't have probes in the same region as the root servers. It's the reason why most of us have Anycast nodes in the 30 to 50-ish range or above for the number of places we're in. I think probably we're in the range of dozens but not hundreds of probes, would be my back-of-envelope guess.

Depending on what's defined with Paul Hoffman's suggestion of geographic regions, if we leave room to say, "We should have a dozen probes per region," and make sure that they are picked to be in topologically well-connected and well-peered locations—something along those lines—I think we might be able to get away with that.

DUANE WESSELS:

These recommendations, in some sense, are all related. I tried to separate them a little bit, in terms of number, location, and then connectivity requirements. But I think we should make recommendations in all of those areas. To me, the first is the number. Warren was talking about that it may be—that the number may depend on the thresholds.

To me, I think that we need to figure out this first. We need to figure out the number and locations first, and then from there, we can do the

---

thresholds. They are obviously related. We have to do one before the other. In my opinion, we should do this one first, and then the thresholds later.

PAUL EBERSMAN: I would tend to agree. Deciding how widely-spread and how detailed our polling topology is, is going to be the harder operational problem. Once we have that, the actual latency that's reasonable is going to fall out of that very quickly.

DUANE WESSELS: Yeah. Some of this, I think, would have to be left up to the organization tasked with operating the infrastructure, but we can make certain recommendations—as we've talked about, geographic or topology-based. There's been a lot of talk about 10 or 12. That's what, Paul, you're using right now in your tests, right? Russ, go ahead. Do you want to ...

PAUL HOFFMAN: Sorry. I was going to answer. I was going to say yes, but that was just out of convenience. It would be just as ... If the work party started to go towards 20, I could set up 20 easily as well.

DUANE WESSELS: Okay, got it. Russ, go ahead.

---

RUSS MUNDY: Yeah, thanks Duane. I was thinking that we could also, perhaps, include in the description of how many probes and the geographic base, because I think since this is really an ICANN and ICANN-tied activity, starting with the five geographic regions makes complete sense. But in terms of the ability to be in a topologically-advantageous spot, there might be a way that we could identify that the probes themselves could check on each other's, and when it was set up, that the delay could be no more than some relatively small number, which would encourage being in a topologically-rich connectivity spot. But I do think that tying to the ICANN geographic regions is something that is, if not totally necessary, I think we'll need to do that to get this sold across the ICANN community.

DUANE WESSELS: Sure, okay. Thinking about numbers, I guess my question is ... It feels like we don't have a lot of technical basis for choosing any numbers. We're just sort of picking a number out of thin air, I guess. Can anyone suggest reasons why, an algorithm, or some process we can follow that would help us settle on a number? Karl, that's probably not the question you were going to ask, but I see your hand is up. Karl, if you're speaking, you're muted. We can't hear you.

KARL REUSS: Sorry about that.

DUANE WESSELS: There you go.

KARL REUSS: Zoom's acting up for me again. Sorry. Yeah, I was speaking about the previous question, but I think it ties in. Many of us have talked about top-end exchange points. That leads us down the path of how do we define top? Is it number of hosts there, or number of people there, traffic? I think it almost becomes easier just to do a regional divide, and then geographic topology inside of that to some degree, because if we try to just pick a single number, how do we divide up where those numbers go?

DUANE WESSELS: Right. I mean, we can certainly pick a number that is divisible by the number of regions. We could say 5, 10, 15, 20, whatever. That seems pretty straightforward. I guess my question is how do we justify 10 versus 100, or do we need to? Russ?

RUSS MUNDY: Yeah. I'm thinking that we've a number of discussions, and we've landed on the point where we say ... Nobody seems to have a strong technical base for saying exactly how many. So, I was going to suggest that we just, at least for this point in time, pick a range—5 to 10 vantage points in each of the five regions. That would be 25 to 50 vantage points.

DUANE WESSELS: Alright, that's a concrete suggestion for sure. I like it. Paul, go ahead—  
Paul Ebersman.

---

PAUL EBERSMAN: This is one where it's not what we do for a living. At my current day job, we were refiguring where to do our Anycast, and we went to our peering guy. We all know folks in the NANOG community. There are certain rules of thumb. If you want to cover most of Western Europe, you go into DE-CIX, AMS-IX, and LINX. There are certain rules of thumb that all of those people can just rattle off like we rattle off EDNS options.

It's probably worth asking those people, because to some degree, what we're really asking is, what is the best transit to the most eyeballs and/or recursive resolvers that serve those eyeballs? It probably follows the traffic of the CDNs and the other folks that have these peering experts.

DUANE WESSELS: This is a little bit related to the work that I did in the appendix here, which is using RIPE Atlas. If you wanted to cover 25% of the root server system is RIPE Atlas probes, you need about 50 or 60. If you wanted to cover 50% of the root server system, you need about 200 or 300. If we're talking 10s, then we're getting about 8 to 10%.

PAUL HOFFMAN: Duane, that's not the root server system. That's every single instance.

DUANE WESSELS: That's right, that's every single instance, yes. Sorry.

---

PAUL EBERSMAN: Right, and it reflects that the RIPE Atlas topography is not terribly even at all, either. It is very Euro-America centric.

DUANE WESSELS: It might, yeah. I don't know. I would say that's a ...

PAUL EBERSMAN: We went through that with ThousandEyes versus RIPE Atlas probes for monitoring our stuff, which had to be worldwide. That's why we started paying money to ThousandEyes.

DUANE WESSELS: Certainly, we know that there are more Atlas probes in Europe, but I'm not convinced that the shape of this graph would be significantly different if the probes were distributed otherwise. Anyway, we're going to run out of time here pretty soon, and I know Warren has his hand up.

WARREN KUMARI: Yeah, I think I'm actually going to say something similar to Paul. I think it would be a fairly bad idea to choose a certain number of probes per region, and scatter them around. For example, if you connect to a network in North America, you're peered with almost everyone in North America, or you can easily reach them over a very small number of hops. If you're in Africa or large chunks of Ocean Asia, you go huge

---

distances to get somewhere, and you go, often, through another continent.

So, I think that the idea of taking regions and saying n per region exposes the belief that there is geographical stuff tied to the network topology, and that it's not network topology tied to network topology. I also think that two or three per general network area would be a much more reasonable number than on the order of 10 per whatever, but that's a separate rant.

DUANE WESSELS:

Are you suggesting two or three per ICANN region? You said 'network area.' Does that mean something different?

WARREN KUMARI:

That kind of means something different. Africa is split, sort of, into two or three rough network areas, whereas North America, I would more classify as one. So, I think that kind of like what Paul said, talking to people who do peering for a living, and saying, "We want to reach this amount of the end letters with about this many nodes. Please go figure that out," would be a much better option. This is more complex. Happy to chat more, but probably not now.

DUANE WESSELS:

Alright. We're running out of time for today. We obviously didn't make much progress here, so we'll maybe try to take that to the list, or save it for the next meeting. The last topic on the agenda was ... Want to talk a little bit about our end goal, which in my mind ... Our next big deadline

---

or milestone is that the RSSAC is having a workshop in October, and it would be really great if the work party could have a lot of these questions answered, and maybe even make some soft recommendations on thresholds, so that RSSAC can cogitate on those and make comments on those at that workshop.

I think before we get to the thresholds, we need to figure out the thing we were just discussing—the recommendations on number and location and so on. We have maybe three or four more of these meetings before the RSSAC workshop. Hopefully, during those meetings, we can come up with a plan for thresholds.

So, I won't keep anyone much longer. Thanks, everyone, for your comments on the transcript and on the chat. If you would like to continue the discussion on the mailing list, please do so, if you have thoughts. That way, everyone can participate. That folks that weren't able to make the call today can also chime in.

Our next call is in exactly two weeks. We have a standing meeting every two weeks. I think Ozan will probably send out a reminder just in case, but hopefully we'll see you all in two more weeks. Thanks, everyone.

UNIDENTIFIED MALE: Thank you.

PAUL HOFFMAN: Thanks. By the way, Duane, I disagree that we didn't make good progress. I think we've whacked off a whole lot there.



DUANE WESSELS: I just meant on the last point. We didn't have a lot of consensus there.  
Yeah, thanks, Paul.

PAUL HOFFMAN: Yeah.

[END OF TRANSCRIPTION]