MONTREAL – RSSAC Caucus Meeting
Tuesday, November 5, 2019 – 15:15 to 16:45 EDT
ICANN66 | Montréal, Canada

FRED BAKER: I'm supposed to do an attendance reminder. What times we're supposed to do that?

BRAD VERD: Ozan collects the attendance here. We keep track. Ozan, go ahead.

FRED BAKER: So Ozan's going to be asking if you're here. If you're not here, don't sign in.

OZAN SAHIN: This is Ozan, for the record. I have attendance sheets which I'll be circulating. I would kind of ask if you could circulate your name. And if you don't see your name, you can just add it to the sheet. Thank you.

FRED BAKER: Okay, great. Thank you, Ozan. Welcome to being here. Let's run around the room and get everybody to introduce themselves. I know the RSSAC people. There might be other people here that don't. So, I'll start with me. I'm Fred Baker, Co-Chair, RSSAC. I'm with a company called Internet Systems Consortium, one of the root operators.

BRAD VERD: Brad Verd, Co-Chair of RSSAC; VeriSign.

UNKNOWN SPEAKER: [inaudible] from Benin [inaudible]. RSSAC Caucus. Thank you.

TOM MIGLIN: Tom Miglin, representing [inaudible] RSSAC; from NASA.

STEVE SHENG: Steve Sheng, ICANN Org, staff supportive of RSSAC and SSAC.

RUSS MUNDY: Russ Mundy, Caucus member and SSAC liaison to the RSSAC and obviously, also an SSAC member.

DUANE WESSELS: I'm Duane Wessels, root zone maintainer, liaison to RSSAC.

PAUL HOFFMAN: Paul Hoffman, ICANN Org, as an RSSAC Caucus member.

WES HARDAKER: Wes Hardaker, University of Southern California's Informational Sciences Institute.

OZAN SAHIN: Ozan Sahin, RSSAC support staff, also managing the remote participation today. And I see in the zoom room we have Jack Biesiadecki, Ray Bellis, Shinta Sato, JPRS [inaudible - 00:02:25] and Yoshiro Yoneya.

HIRO HOTTA: Hiro Hotta from WIDE + JPRS; RSSAC member.

MICHAEL CASADEVALL: Michael Casadevall, RSSAC Caucus member.

KEN RENARD: Ken Renard, Army Research Lab; RSSAC.

KEVIN WRIGHT: Kevin Wright, RSSAC member from DISA.

KARL REUSS: Karl Reuss, University of Maryland; RSSAC member.

DAVE LAWRENCE: Dave Lawrence, Oracle and RSSAC Caucus.

NAVEED BIN RAIS: Naveed Bin Rais, RSSAC Caucus.

MOHIT BATRA: Mohit Batra, I'm working as a consultant with the Indian IT Ministry, RSSAC Caucus member and an ICANN Fellow. Thank you.

DANIEL MIGAULT: Daniel Migault, IAB liaison to RSSAC, Erickson. And I'm here as a RSSAC Caucus.

LARS-JOHAN LIMAN: Lars-Johan Liman, Netnod; RSSAC member.

ADIEL AKPLOGAN: Adiel Akplogan, ICANN Org and RSSAC Caucus member.

JEFF OSBORN: Jeff Osborn with ISC; RSACC member.

UNKNOWN SPEAKER: [inaudible], just an observer.

YOSHIRO YONEYA: Yoshiro Yoneya from JPRS, observer.

SHINTA SATO: Shinta Sato, JPRS; RSSAC Caucus.

UNKNOWN SPEAKER: [inaudible] from JPRS, observer.

UNKNOWN SPEAKER: [CROSSTALK], support staff.

FRED BAKER: Okay, thank you. The agenda that we have today is actually the same agenda as the Caucus will have on the Sunday at the IETF meeting in Singapore, and it's in front of you. So we're going to talk a little bit about RSSAC Caucus engagement; when do we meet, and that kind of thing; talk about some of the recent publications that RSSAC has produced, and current work that we have going on in work parties in the Caucus, and then some other work that we have really just started up in the RSSAC. And then we'll take any other business.

Question: does anybody have a change that they would like to make to the agenda? Do we have a topic to throw in AOB? Hearing none, okay.

So Caucus meets at every ICANN meeting. This is an example. It meets every other -- the even-numbered IETF meetings -- and so, we will have a meeting in Singapore in two weeks and -- well, a week and a half. And as I said, it will actually have the same agenda. If you contribute to a work product that is being done in the Caucus, your name will be on there as an observer or a contributor, depending on what you did.

And I want to make that point fairly clearly because, frankly, we've got a lot of people in the Caucus that don't contribute, and we're concerned about that. So we're trying to figure out who's actually doing work in the Caucus and do some work there. So, encouragement to you, please feel free to get involved in the work that's in the Caucus.

Recent publications, we've actually done a few in the last six months. The most recent one was the workshop report. The RSSAC had a workshop, roughly a month ago, in Reston, Virginia, hosted by VeriSign, and primarily working on the metric stuff. We also talked briefly about the resolver work, but primarily, dedicated to the metrics effort. And I thought it was actually a pretty productive meeting. A lot of good things happened there.

Current work parties, we have two. One of which is RSS metrics, the fundamental question being, what's good. How do you measure the RSOs, how do you measure the RSS and determine how well they're doing? And that would be in terms of actually delivering the IANA data set, latency, and that kind of thing. The chairs are sitting over there, Duane and Russ. And Brad, we've got your name on this. I'm not entirely sure why, but --

BRAD VERD: I will defer to Duane.

FRED BAKER: I thought you might. So, Duane and Russ, do you have anything you want to say about the metrics party right now?

DUANE WESSELS: Well, certainly at this meeting, if we're just going through the agenda, then no. But I have some slides prepared to update people on the status of the work party.

FRED BAKER: Thank you. And when it comes to the resolver work party, I'm going to call on you, Paul. Paul's been doing yeoman work there.

PAUL HOFFMAN: This is Paul Hoffman. We are basically done with the software for the resolver work. And my last to-do item was to create a page -- it's essentially two pieces and such like that -- and I was about to do that and someone reminded me that -- because it has intentionally an open license, anyone can use the software. And yet, I had not cleared that with my lawyers. So, that will happen soon. They insist that they work for us, they just want to make sure how to do it right.

But I can't really go and ask them for forgiveness afterwards. I believe I will have that answer in a week and a half before we have the meeting in Singapore. So that would close out the work party work.

Basically, it's two pieces of software. And we haven't had much interest in using it yet, but as soon as someone is, then we might reinitialize. Because the output of using that software could be reports and such like that. So the software's there and it's still the same URLs, its just I need to make sure I actually know what the license is.

FRED BAKER: Right. Then we have some publication updates. And we'll go to AOB after that. While I was reviewing the agenda, we've done certain parts of it. So Duane, you want to show us some slides?

DUANE WESSELS: I'll wait for Ozan to put the slides up. So the plan for this meeting is -- This is going to be sort of a high level status update on the metrics work party and our documents. This is similar to a presentation that we gave at the start of the RSSAC workshop a month ago.

In fact, it's sort of a copy and paste from that, plus updates from the last couple of weeks. So as we're going through this, if you see something that doesn't look right, please ask questions or point it out. It's likely a mistake on my part from updating the slides.

So this shows the high level structure of the metrics document, as it exists currently. There are these, I think nine -- actually it's maybe not, maybe it's even changed a little bit recently -- but there's about nine sections.

So there's an introduction, there's a background and scope, there's a whole section devoted to vantage points, and then there's a relatively long section about general points about metrics and measurements. We get into the some of the specifics about metrics in the RSI metrics, and you might notice here, this new three-letter abbreviation, RSI. This is Roots Over Identity.

So the document has just recently been changed from using RSO when it really meant RSI, roots over identity. So, we'll try to use that going forward. Then there's the RSS metrics section. We have some recommendations, and a section of example results. Until just recently, there was a section dedicated to possible future work. I believe this is a little bit flux at the moment. It might end up as a recommendation or it may stay as a separate section. We're not sure yet. So next slide, please.

So the introduction is relatively straightforward. And this hasn't really changed over time. So there's probably not a lot to talk about here. But it has a statement of work. It, again, goes over the structure of the document as a whole, and has some boilerplate about being an RSSAC publication and so on. That's relatively straightforward. Next slide.

There's a there's a background and scope section. This is evolved a little bit, I would say, in recent months, with some additions. In general terms, there's a couple of paragraphs to talk about the purpose of the metrics and how they're focused on minimum levels of performance. We did have some discussions a month or two ago about whether or not we should also be talking about good levels of performance in addition to minimum levels.

So we've settled on only talking about minimum levels of performance, and any work on good levels of performance that would that would be different, is reserved for possible future work. The purpose mentions RSSAC037 as an impetus for doing this work, but not necessarily. The only reason for doing the work, and it does mention some of the terminology from that 037 document, such as PMMF, that's likely to have a different name when it comes out of the governance working group work that's about to start. And in this purpose section, we also specifically mentioned SLEs as service level expectations.

Now, there was a time where I could say that the document didn't ever mentioned SLAs. But that's no longer true. I think we mentioned SLAs is as a possible future work item. It doesn't have a lot to say about SLA, but that that term is in there now.

So we have a section that talks about uses not in scope for this document. So this was good that we nailed this down a few months ago, to help focus our work here. So things that are out of scope include, research into performance trends of the root servers. That's not why we're doing these metrics. And they're not also to be used for making comparisons between roots of our identities. And you'll see one of the ways that we accomplish that, later on we'll talk about how for the RSIs, we're only publishing pass fail metrics.

There's a section on prior work, which references a lot of existing RSSAC documents and ways that it's related to this work or maybe even not related to this work. And then there's a terminology section. A while ago, the terminology section had some entries which were actually slightly different than the RSSAC 026 terminology document.

But those differences are sort of in the process of being resolved. And we're going to be consistent among all of our documents. And so the terminology got a little bit shorter, until yesterday, when I actually added a couple of more added vantage point and collection system to the terminology. So take a look at that if you're so inclined. Next, please.

These are the five terms that are sort of defined, specifically, in this document. We have measurements, which is sort of a very small unit. It's like a single query response at a given interval. And we have then metric, which is sort of a way of aggregating all the measurements together to get a longer term result. And we've settled on monthly metrics. So all the aggregation intervals will cover a period of a month. We have a definition of threshold, at least, how it's using this document. And then, as I said, I've added vantage point and collection system. Next.

Okay, so there's a vantage points entire section and here, the document says that we recommend that approximately 20 vantage points be deployed for the purposes of these measurements and metrics. The location says they should be distributed evenly among five geographic regions. I think now it maybe even says something like, "distributed approximately evenly among five geographic regions.” And we have a future work item, which would make improvements to the distribution of the probes at some at some time in the future.

There's a little bit about connectivity, which sort of stipulates that the vantage points should be located at data centers with reliable power and good connectivity. And lastly, vantage points within the same broad geographic region should use different connectivity providers, if at all possible. So that's designed to get good coverage and not get locked into certain networks.

There's a section on sort of general topics, and it's actually becoming quite long. There's been a recent suggestion that maybe this should be split up a little bit, and we can look at that. But for now, there's these 10 sub sections here. There's a section on reporting, which as I said, we've agreed that the reporting for these metrics should be done on a monthly basis.

Long time ago, or maybe not too long ago, we had daily on there, so I think this is a good change. Monthly is better. There's talks about how timestamps should be represented and how measurements should be scheduled within their five minute intervals and so on.

There's a bit about lapsed time, how you calculate lapsed time between queries and responses, especially in the context of TCP and how you deal with timeouts. Connections errors, we have a little bit about how certain connection errors are treated as invalid data or timeout, and those measurements get ignored.

The section on spoofing is to make the point that the implementation of this measurement system should do all the standard things that DNS software has to do to protect itself from spoofing. So that's making sure that the answer section matches the query that you sent and that the query IDs match, and that the UDP port numbers match and that they're randomized, and all those sorts of things.

The Anycast section says that the measurements are not -- they're directed at the Anycast service addresses and not at any specific instances of a root server operator, so that the routing system does its job of routing packets to the network and the queries land where they would naturally land for anyone else. In these measurements, we're not trying to uncover or explore or make any statements about any past deployment of the operators.

The note about measurement reuse is in reference to the fact that, for some of the the RSI measurements, essentially, we make one query and use it in three different places. So we use it for availability, latency and the publication latency. Unexpected results talks about how the measurement system should deal with the results that are errors or otherwise unexpected. They should be recorded and it may be necessary to refer to them later, in discovering certain problems, especially like in the case of correctness checks.

Getting down to the end here, one of the new sections, which I have a whole other slide about, is determining the number of our size required for reliable operation of the RSS. This was something we spent a lot of time on last month in our workshop. And this last section is actually quite new. Its potential effects of metrics on independence and diversity.

And so, the point here is that, to reiterate, that the root servers, the root server operators value, their independence and diversity. These metrics are not intended to change the way that they deploy their systems and their networks, but rather just to keep them to those minimum service levels. So if you haven't had a chance to look at the last two, I would encourage you to do so.

Okay, so as I said, we spent a lot of time on this before, talking about all sorts of things -- availability and whatnot -- and we agreed that one of the first things we had to decide was settling on a number, the number of root servers that had to be available for reliable operation of the system as a whole. And the formula that we came up with is shown here.

It's essentially two thirds of the number of operators minus one, rounded up. So what this means is that, if there's a DNS client that's making a query, and the first one gets a timeout, then its second query would be successful with two thirds probability. So for N equals 13 root server identities, this results in a value of K equals eight.

And the document includes this graph to make the point that the value of K depends on the number of N. So if the number of root server identities changes in the future, then the value of K also would maybe change, maybe not; depends. Due to the rounding up, maybe, maybe not.

WES HARDAKER: A quick question, Duane.

DUANE WESSELS: Yes.

WES HARDAKER: So up means -- above up, is generally fine. We don't need to worry about it. Is there a notion that just below eight, say seven, is not necessarily fully broken down and down? Is there text in there about that? I'm sorry for not scrolling really quick to find out.

DUANE WESSELS: That's explained in the section about the the RSS availability metric, so that's a little bit later. But the reason this is earlier is because we also are using this concept to determine the availability threshold for an individual identity.

WES HARDAKER: Okay. Yeah, that's fine. It was just, at one point in the past, I think we were more than five down, under a major attack.

DUANE WESSELS: Yeah, that did change. My initial proposal from a month ago was more like a step function, either you're all good or you're all bad. But that did get changed. So now it's a smoother --

WES HARDAKER: Right. I don't think the world noticed. That's sort of my summary point.

DUANE WESSELS: Okay, lsure. Let's go to the next slide. All right. And you notice a mistake here, because there should only be four bullets. This section is the metrics on the root server identities. There is availability, response latency, At one point, we had two separate correctness metrics. And those have now been sort of merged. There is just a single correctness metric, which I will discuss. And there's the publication latency metric.

So, measuring the availability of a root server identity is relatively straightforward. The measurements performed separately over all the different transport protocols of V4, V6 and UDP TCP. And over the aggregation interval, you send some number of Q queries and receive some number of R responses, and then the availability is just the ratio of those two. And what we've agreed on as output of the workshop was that the threshold for this metric should be 96 percent. And the next slide explains why it's 96 percent.

So we apply this simple parallel model K out of N availability model to the root server system. And if you do some searching for that phrase, you'll find papers and articles that that talk about this formula. This is a formula for calculating overall system availability, given some number of components, N in which K of them need to be up at the same time.

This is what I call simple parallel K out of N availability, because it makes certain assumptions about your system. It assumes that, for example, all the components are identical and independent, which is not necessarily true for the root server system. But we agreed to use this model anyway, rather than deal with the very much increased complexity of a more sophisticated model.

So in addition to agreeing that we would set K equal to eight for the current system, we also agreed that it's desirable to have an overall system availability of five nines 99.999 percent. And when you plug in these numbers of 99.99 percent, and N equal 13 and K equals eight, you end up with an overall identity availability of 96 percent. That's the value of little A in this formula. So that's how we came to reach 96 percent.

The next metric is the root server identity response latency. And, as with the previous one, this is measured separately, over all the transports. It actually reuses the same query and response from the previous one. And the calculation of this is relatively straightforward. You just aggregate all of the latency values from all the probes over the one month period, and calculate the median.

And if you had a measurement from every probe, from every server at every interval, you would have a maximum of -- in a month with 30 days -- you would have a maximum of 172,800 possible measurements that you would aggregate. The thresholds for this metric, keeping in mind this is a medium that thresholds -- we initially agreed on 250 milliseconds for UDP and 500 milliseconds for TCP. Following the last work party meeting, there was some discussion that maybe this should be raised up to 400 and 800.

We still haven't fully resolved this, but we're working on it. There's still some homework for us to do, or at least for me to do, and analyzing some RIPE ATLAS data to see what things currently look like. But I'm hoping that we can get agreement on 250 and 500. Go to the next slide.

So this is kind of what I was just referring to. This is a latency distribution graph from the RIPE ATLAS system. This is from the anchors only. The anchors are like One.IU boxes that generally go in data centers. They'll have more computing power and they can do more measurements. They tend to be a little bit better connected.

This is from all of the RIPE ATLAS anchors, which, there's about 500. And this is for the month of September. And in my data, I have, essentially, one measurement per anchor, per day. Which is, of course, a little bit different than what we're suggesting here where you would have one measurement every five minutes. So this is for one of the root servers. And you can see the two lines for IPv4 and IPv6. And this is a cumulative distribution function.

So what we're suggesting, or what we're proposing, is that the threshold is against the median. So in this case, the median value would be halfway up along the Y axis, where those lines are. Somewhere in the range of, looks like 30 to 40 milliseconds, would be the median values for this particular root server. And then that vertical line off to the right a little bit, that's the 250 milliseconds threshold. So this is just one server.

We do have the data for some of the others. But also, the future work that I intend to do on this, is to do an analysis a little more similar to what's being proposed in this metrics, which is to have, instead of 500 vantage points, have closer to 20, and see what that looks like.

Okay, the last -- nope, not the last. The second to the last, root server identity metric, is correctness. And I'd like to thank very much, Paul, for doing a lot of this work. This has kind of been where he spends most of his time in the document working on this. And if you go and read it, you'll see that it's actually quite detailed. There's a lot of stuff there. So, the correctness is based on both what you could call exact matching, as well as DNSSEC validation.

The exact matching works by comparing the data in a DNS response message to a copy of the data from a recent root zone file. So the collection system has to keep copies of root zone files so that it can do those comparisons. This particular metric is not measured separately for each transport. The transport is just chosen at random for a given interval. So there's only one metric out of this.

The way this works is that 90 percent of the time you would do a measurement, expecting to get a positive answer for a data that does exist in the zone. And 10 percent of the time, you would do a query that you would expect to result in NX domain. And the document says what the query name should look like, it's shown here.

And I wanted to highlight that because what we came up with was to put the name of this document, RSSAC 0XX -- whatever that ends up being -- replace the XX with actual numbers so that if people happen to look at some of this data they'll know where this is coming from. The last component of the query name is of course random, so that it's going to be an NX domain.

DAVE LAWRENCE: Duane, a quick question. I'm sorry, I don't remember this part of the document right now, this particular detail, but random could collide possibly, because it's random. And if it did, it could possibly also collide with a domain that is legitimately, say, wildcarding, because that's the domain policy, to wildcard their effects.

DUANE WESSELS: Yes. So there are some details being left out. I think, if I remember correctly, Paul, it stipulates a 10-character random string, so less likely to collide with short gTLDs. But yes, there could be a collision, in fact, you could get back a positive answer when you expected a negative answer. And the checking is sort of designed to account for that.

Like, if you get back a positive answer, you apply the positive rule checking, not based on what you expected to get. So even if you did randomly choose com or whatever, you would you still be able to verify that it's the right result, the right response.

DAVE LAWRENCE: I'm sorry. I think I just forgot to identify myself. So for the record, since we do that here, this is Dave Lawrence, Oracle.

FRED BAKER: So, I have a question for you, Duane. I'm looking at this fraction, and it occurs to me that there might be -- it could be weird. Let's suppose that I send queries to something all day long. I get one response, and it's wrong. Or no, I get one response, and it's correct. So is my correctness 100 percent? I'd like to believe that the denominator had to do with how many requests were sent, not how many responses I get.

DUANE WESSELS: Well, if you do it that way, then you have a complication with timeouts and things like that. If you sent 100,000 queries and you got 99,999, you're saying that's not 100 percent correct?

WES HARDAKER: Let me put it another way, Fred. Yes, you're correct. Something else is going to be horribly wrong. So this metric isn't designed to catch that case. Another metric would catch the fact that 99 percent of your requests weren't being answered.

FRED BAKER: Okay, as long as it's caught.

DUANE WESSELS: And additionally, something that's not representative of the slides, is that what we are recommending -- we have some sample reporting, and when you report the results, you would include the denominator. You would say, this correctness is based on one response or a hundred thousand, or whatever. So you would be able to see those cases where something went obviously wrong and you got way fewer measurements than you expected, and then you could investigate.

So, the threshold for this metric is 100 percent. That's been something that everyone has agreed on from the start. So, that's great to see. Next, please. The last roots over identity metric is publication latency. This reuses the queries from the availability metrics, so that's good.

Since the measurement intervals every five minutes, the resolution of this measurement is limited to five minutes as well.

This metric is a little bit complicated because something, essentially this collection system, needs to know the time at which new root zones got published, and then do the calculation of when it sees a new serial number from each of the servers at each of the vantage points. So there's a little bit of complexity there, but I think our description is good and understandable. The latency against which is measured, is the median latency of all the vantage point measurements for each serial number per day.

Here it says all serials per day, I think that's actually incorrect. Oh, no that is correct. I'm sorry. It's one result per -- but it should be month and not day. I think that's an error. Sorry. Yeah, this should say month. So for each month, there is just one publication latency result. What this does mean however, is that if there, for whatever reason, are our days where there's more root zone serial numbers published than others, then that increases the total number of measurements that go into the aggregation here. The proposed threshold for this is 65 minutes, that's twice the SOA refresh value, plus one five minute interval for boundary conditions and whatnot.

PAUL HOFFMAN: Duane, the five minutes got added in the last meeting. I truly did not understand. I don't think you were the one who brought it up, but I don't remember who brought up the -- and it needs to have an extra five minutes. I think that's actually antisensical.

WES HARDAKER: That was me. So the hypothesis was that we were trying to catch the ability -- give somebody one additional time period in order to get it wrong. And because your measurement will coincide, possibly exactly, with the point at which somebody starts pulling the zone -- and it takes more than a second to pull the zone -- that if you push it to plus five minutes, we will give them the latency that they want the whole time.

PAUL HOFFMAN: This is Paul Hoffman again. That would be true for an individual measurement. The threshold is based on the entire month's worth of measurements. And so, I would imagine it would be exactly two, still; That you don't need the five minutes unless you wanted them to be able to miss two SOs and a little bit, for every single time in it, which to me, doesn't make sense.

WES HARDAKER: Yeah. And I remember we were talking about it the other day. We were talking about it in this section and I think we said we were going to do it to both, and I agree I'm not sure it's as necessary for this one. It could go either way, though.

DUANE WESSELS: Alright, we'll mark that for further discussion at a later time. All right, so that's it for the four RSI metrics. So next, we have the corresponding for RSS metrics. This one's a little bit a little bit tricky, and I think this gets to -- whose question was it? This was your question, Ws, about how do we deal with this. So, in each measurement interval at each vantage point, it should expect responses from at least K identities, so at least eight identities. That's what the term R t,v represents.

So, you could count how many root servers responded to the SOA query in each interval at each vantage point, over all those intervals, you calculate the sum of either the minimum value of K or of R. So if all the R values are greater than or equal to eight, then this is just a sum of eights for how many intervals you have. And then you divide that by the sum of K over those same intervals.

So, if in every interval the R values are eight or greater, then A is 100 percent. If you have one interval where R is a little bit less than K, then it starts to go down. We start to see a decrease in availability. And I have a table that has that example and some other ones, on the next slide. So the threshold for this is five nines. Does the formula make sense, in this brief explanation?

WES HARDAKER: I understand the formula. I have not come to conclusion about a best yet.

DUANE WESSELS: Okay. So let's take a look at the table, which has somewhat contrived examples. And I apologize, this is a little bit hard to read. It's a copy and paste from the Google doc. So here's some examples of hypothetical situations and how they would impact this availability metric. So our reporting interval is one month, so that's what all these are sort of based on.

If you have a month-long attack that manages to take out a single root server identity entirely, the availability is 100 percent. Because in all the intervals, the value of R is 12, which is greater than eight. If you have a month-long attack that takes out five identities entirely, you still have 100 percent availability, because again, in all the intervals, R is equal to eight, which is equal to K.

However, if there's this magical month-long attack that takes out six identities entirely, then the availability is at 87.5 percent or seven eights, because in all intervals, the value of R would be would be seven. If there is a 24-hour attack that takes out all servers entirely, the availability is 96.66 percent. Where it says 2930 in that box, that's supposed to be a fraction. That's 29 divided by 30. So that's assuming a month with 30 days. And in one of those days, all Rs equal to zero. So you end up with a fraction of 29 over 30.

If in one five-minute interval, one vantage point can only reach seven identities, but in all the other intervals it can reach all the identities, then the availability is 99.99992 percent. So that's six nines. The point at which you get close to the five nines availability threshold is when you have about 14 such intervals or if you have, say two intervals where seven vantage points can reach no root servers, then you're at about five nines, 99.9989 percent.

So, it has been pointed out in the document that some of these-- these are contrived cases, and in some of these, they may be unrealistic. Because if you've got a vantage point that that can't reach any root servers, then it's probably a vantage point problem, it's not a root server problem. And so, we handle that in other ways.

There is text in the document that says the vantage points have to do connectivity test to make sure that they're online and connected to the internet. And if you have a case where it's offline, then you discard those measurements so they're not included in this.

WES HARDAKER: If it makes you feel better, when I was doing what's the minimum amount of time to -- I did the math in an opposite direction and came up, what's the most that you'd have to be down in order to fail, and five minutes was about right. So that's sort of matches what you're saying.

DUANE WESSELS: Any other discussions about this? This is something we talked a lot about in our workshop in Reston.

BRAD VERD: Well, this was a topic brought up by your counterpart who was representing you -- channeling you there at the--

DUANE WESSELS: All right, we can move on then. The RSS response latency. So in this metric, we also have agreed -- this is a case where we use this value of K -- so for this metric, and from each vantage point, you find the K lowest latencies and then you aggregate that subset of the -- calculate the median value of that subset of lowest latencies. And since it's a subset of sort of the best, the thresholds here are lower than they were for the the RSI case. Thresholds are 150 milliseconds for UDP and 300 milliseconds for TCP. All right, next.

So there is an RSS correctness metric. This is very straightforward. It's just a simple aggregation of all the measurements from all the identities, using the same ratio, the correctness, the number of correct responses divided by the total number of responses, and the threshold is 100 percent.

And something that we've talked about a number of times in the work party and other places are that, yes, this is a high threshold, so any time where the metric is not met, doesn't meet the threshold, there needs to be a good understanding of what's really going on here.

Did this response really come from a root server or was it some attack [inaudible] or was it some measurement failure or something else. So, even though this seems like a very high bar, I think we've given ourselves lots of chances to explain any anomalies, revise things if necessary, and so on. Okay, next.

Similarly, the RSS publication latency is just a very straightforward aggregation of all the measurements from the root server identities. We calculate the median of those publication latencies. And the threshold here is 35 minutes. That's one SOA refresh interval, plus the five minutes wiggle room that we sort of talked about.

And we talked about this just yesterday, the day before, explaining why -- and I think the document now explains why, but the reason for this being 35 and the other one being 65 is that you wouldn't really expect that -- in order for this threshold to reasonably set at 65 minutes, you would have to assume there are cases where all of the root server identities are sort of right at that threshold all of the time. If there are cases where individual root server identities are above 35 minutes, you wouldn't expect all of them to be there. Next please.

So as I said, the proposed reporting for this is that for the RSI metrics, the report just has a pass-fail indication. However for the RSS metrics, the actual numbers will be included in the report. But in both cases, raw data will be available for anyone who wants to see it. That's one of the recommendations, that the raw data must be available. But the high level reports will not have the actual numbers for the RSI metrics. Okay, next.

This is an abbreviated sense of some of the recommendations in the document. And this has been changing recently so I forget exactly which recommendation number it is, but there is a recommendation that the ICANN board commission a proof of concepts implementation of this measurement system.

And then there's another recommendation that places some requirements on what we call the official version of the measurements, that must meet the requirements regarding the number and location and connectivity of vantage points that were specified here. The software used to do this must be published as open source.

As I said, the raw measurement data available to anyone, in the interest of transparency. Only pass-fail indicators for the reports for each root server identity, and that the raw measurements would need to be shared with root server operators in cases where the thresholds were not met. So this was changing recently. Did I miss any other recommendations that we--? We had a third recommendation, did I not capture it?

STEVE SHENG: Yeah, I think the third recommendation's about future work.

DUANE WESSELS: I don't think I put that in here, but okay. You can go to the next slide, please Ozan. There's a section of example results, and it sort of looks like this. I just changed just the other night, because initially, there were thresholds in here that were intentionally silly, so that when we were looking at these we didn't get focused on these being the proposed thresholds.

But now that we're past that, the thresholds in the example results are the same as the recommended thresholds for all the metrics. And there's lots more of those, so if you want to go look at them, feel free. This is just the first one. And that's the end of the slide deck that I had. Let me know if I missed anything important.

NAVEED BIN RAIS: Naveed for the record. If you can just go back to the graph that shows the future expansion of the root versus the K factor. What I got from that is that you're using K as seven there, rather than eight. So I just wanted to understand that maybe. I'm not sure. Because the eight gets as soon as 11, for example.

So it means that from number of root servers 11, you get an eight value, right? So in that case, if three of them fail or more than three failed, then you don't have the same kind of threshold or the availability that we are expecting. So I just want to understand this point.

DUANE WESSELS: Okay, sure. So the reason that I wanted to include this graph in the word party document, is to be clear that if the number of roots server identities changes, then the value of K also has to change.

NAVEED BIN RAIS: That, I understand. But for 13 here, it is showing nine, rather than eight, which we understood previously. So that may be because your ceiling -- so I'm not sure which one actually we are going to use. So if you see this graph--

DUANE WESSELS: You're saying the graph is wrong?

NAVEED BIN RAIS: Yeah, there might be some value. I'm not sure whether it is intentional or not. So I'm just checking. So for 13, it is giving us nine. So, yeah, maybe.

DUANE WESSELS: Well, the reason I say maybe is because the X axis numbers don't line up directly into the bars, but yeah, I see what you mean.

NAVEED BIN RAIS: The second one is for the response that you get, the values that we are choosing in the case of RSI as to 250 and 500. And you said it's still ongoing and 150 and 300 for RSS. Why this double ratio between UDP and TCP? Iis there like some measurement behind this? For example, for TCP, do we assume that only one query per TCP connection? What if there are more than one? So in that case, that would come down as a median, I'm talking about.

DUANE WESSELS: So, the reason for doubling it is to allow for the fact that there's a connection setup delay with TCP, so that that setup time is included in the measurement. I don't know if the document says only do one measurement per connection, but it probably should. Because really, given that the measurements are over every five minutes, that's too long really, to expect a TCP connection to stay up.

So it has always been my intention or assumption that there would be one query per TCP connection. There are also, as I mentioned in the document about TCP fast open, and that it should not be used, if possible, because that can sort of confuse the latency calculations. So, I guess we need to add something about one measurement per connection.

HIRO HOTTA: About this graph, the left most proper bar is for any code to write? So, we should move the 5, 10, 15 20 just a little bit left.

DUANE WESSELS: Sounds like I have volunteers for making a better graph.

HIRO HOTTA: If we have n 13, we have k eight.

DUANE WESSELS: Okay. Kevin?

KEVIN WRIGHT: Your last slide will kind of visualize my question, And it's about the combinations of the different protocols and measuring the four different combinations. Does it make sense to weight the certain combinations differently, since we get a lot more say, IPv4 UDP, or is it just simplify it and make them all equal weight?

DUANE WESSELS: I'm not sure what you mean by weight them exactly. Certainly, you could give them different thresholds if you wanted to. The only time we do that is in the case of the response latency metric. But right now, each one stands alone. They're not combined together in a way that you could weight them.

KEVIN WRIGHT: For your RSI latency, don't you combine the four combinations into one? Or are all four separate?

DUANE WESSELS: Right. So what we're proposing here is that they're separate. And that's why there's these four rows here. So you could, in theory, pass on V four UDP and fail on V six UDP, or something like that.

RUSS MUNDY: This is Russ. I was going to comment in response to the graph that shows the K over N number, and I believe that was put in originally just to give us a scaling kind of illustration. And it's something that if the work party thinks it needs to stay in -- it does need to get adjustments -- but it can also just simply be eliminated from the document and just based purely on the formula.

DUANE WESSELS: I apologize for getting this wrong. I remember, this was made with [inaudible] and I thought I was very careful with my script. But my feeling is that it would be good to keep this in the document, because I don't want people to think that K equals eight forever. I want people to understand that it's dependent on other factors. And that it might equal eight today, but it might not always, in the future. For that reason, I would like to keep it, after fixing it.

FRED BAKER: Speaking for myself, nothing more, nothing less, I think the graph is useful. But yes, it needs to be correct. Otherwise, we'll have this conversation every time.

KEN RENARD: While you're at it, should that be number of RSIs at the bottom, as well?

DUANE WESSELS: It should. If I was allowed to run gnuplot on my work laptop, it would have been fixed. But it's on my computer at home, which I can't access right now. So I didn't fix it. I'm not sure where we are on time, but I think we took up quite a lot of time, so we should probably be done.

FRED BAKER: Well, we do have one other topic to discuss.

DUANE WESSELS: I think we're done.

FRED BAKER: I had one question for the two of you, before we go on. And that is, what do you see the process on this -- what we've been having is every other week, calls and Caucus and RSSAC, everybody potentially involved. And you've done yeoman work, for which, I thank you. And we've now arrived at a point, it looks like, where we're mostly done. what's the prognosis? What do you have going forward?

DUANE WESSELS: I think we're in agreement that it's nearly done. I would like us to be able to produce a final draft, which would go out to the Caucus, as a whole, for its review in entirety. And I'm sure there will be things found and things to fix, but while we do that, I would ask that Caucus -- this is no longer the time to propose big additions, this is for fixing little things. And then after that review, then bring it to the RSSAC for final. I think, maybe we're a couple months out from being totally done, something like that.

STEVE SHENG: Thanks, Duane. I think in practice, the Caucus review usually takes about a month. That's usually the time we gave to Caucus to review other documents. Following what Duane said, the RSSAC made changes that all the working group discussions happen on the Caucus.

So when a Caucus sees this document, this is hopefully not the first time they see it. All those mails going for us, every work meetings, documents, they already have this draft for a while. So we think, probably one month is enough, more than enough for that purpose.

FRED BAKER: Well, that's all true, except we made changes in it this morning. We made changes in it during the workshop. If I was just a random Caucus member, I might not have noticed. So, I think we're at a point where we need to catch up with them. So, you think you might have that reviewable draft by Singapore, by December? When?

DUANE WESSELS: Yeah, I think, by the Singapore Caucus meeting, for sure.

FRED BAKER: And so, then we can start a month clock on that, at that time? Okay, that sounds good. We have one additional thing on the agenda, and that is that we have three documents that we're revising. Two of those came out of your work party. And those are the RSSAC 23, history, and 26, which is lexicon. And there's some work going on there. And I don't see Andrew here. I wish I did. Steve, can you comment on that? Oh, and the third thing is an update to RSSAC2, which is statistics. So, Steve.

STEVE SHENG: Sure. Andrew is at the SSAC meeting, so I'll stand in for him. The RSSAC made a decision to revise the three documents: the RSSAC 002, that's the measurement, 23 is the history document, and 26 is the terminology.

I think with the 26, the goal is to align the terminology with whatever the terminology come up with this document. And I think Paul is the shepherd for that, as well as the RSSAC 23. The staff is going to do some work to pull the contents together for the Caucus to review. So I think that's just a quick update on that.

FRED BAKER: Okay, thank you. Yes, Ken?

KEN RENARD: I could say a few things about 002 if you like. I just posted in the chat, the link to the current working document for the updates to RSSAC002. This is a Caucus work effort, not a work party, and it's different from the metrics work party. Just invite everyone to review the document, review the changes. And just knowing that RSSAC2 is self-reporting metrics for the purpose of determining trends and performance of the root server system. It's not what we just talked about in the work work party.

But a few things, we did clean up some terminology, some data formatting. There's a proposal to put in measurements for Q name, but we're really trying to keep the scope here. Not putting any research topics in for measurement. Of note was removal of counting unique number of individual IPv6 addresses, kind of putting that down to just slash 64s and unique 64s. So, I invite the Caucus to look at the document, look at the edits. Please comment in the document or via email to Caucus list. And, thank you.

FRED BAKER: Okay, great, thank you. Now we're at AOB. Did anybody have anything to talk about? I don't know of any topics. Okay, so our next meeting will in fact be in Singapore, and similar agenda. We'll be on Zoom for that, so for people that don't attend, they can be on Zoom. That meeting will be on Sunday the 17th, and it's at 3:30 Singapore time. So 3:30 to 5:00. I think with that, we're done. Okay, so let's adjourn.

BRAD VERD: Thank you, all.

[END OF TRANSCRIPTION]