FRED BAKER: Okay. So, the time is now 3:30. Time for us to get started on this. I can hear me. So, it's 3:31, actually. Time for us to start the meeting. This is the RSSAC Caucus group. If you think you're somewhere else, you're in the wrong place.

The agenda is on the screen. Does anybody want to change the agenda in any way? Seeing none. Okay.

Attendance reminder. So, we're asking people to show up, right?

ROD RASMUSSEN: And people to sign in.

FRED BAKER: Oh, okay. So, Ozan is running around with a sign-in sheet. Please do so.

OZAN SAHIN: Fred, you will also do introductions. So, I will circulate the microphone.

FRED BAKER: Okay. Introductions of who?

ONDREJ SURY: Hi, this is Ondrej Sury from ISC.

UNIDENTIFIED MALE: [inaudible], ICANN.

RAY BELLIS: Ray Bellis, ISC.

TERRY MANDERSON: Terry Manderson, ICANN, root server operations.

DARREN KARA: Darren Kara, ICANN.

RYAN STEPHENSON: Ryan Stephenson, DISA.

YOSHITIKA AHAREN: Yoshitika Aharen from JPRS.

SHINTA SATO: Shinta Sato, JPRS.

KAZUNORI FUJIWARA: Kazunori Fujiwara, JPRS.

AKIRA KATO: Akira Kato, WIDE Project.

ROBERT STORY: Robert Story, USC ISI.

KIM DAVIES: Kim Davies, IANA.

KEN RENARD: Ken Renard, Army Research Lab.

BRAD VERD: Brad Verd, Verisign.

DUANE WESSELS: Duane Wessels from Verisign.

RUSS MUNDY: Russ Mundy, Parsons.

PAUL MUCHENE: Paul Muchene, Georgia Tech.

MATT WEINBERG: Matt Weinberg, Verisign.

RICH THORNBERG: Rich Thornberg, Verisign.

SHUMON HUQUE: I'm Shumon Huque, Salesforce.

IHTISHAM KHALID: Ihtisham Khalid, FAST University, Pakistan.

PAUL HOFFMAN: Paul Hoffman, ICANN, OCTO.

HIRO HOTTA: Hiro Hotta, JPRS.

WILEM TOOROP: Wilem Toorop, NLnet Labs.

OZAN SAHIN: And this is Ozan, ICANN staff.

MARY WONG: I get to do that, too? Mary Wong, ICANN policy staff.

FRED BAKER: Okay. So, Ozan, do you want to introduce yourself?

OZAN SAHIN: I already done, but Ozan, ICANN staff in support of RSSAC.

FRED BAKER: Okay. And I'm Fred Baker, ISC. Actually, most of the people in the room were late to one RSO or another. We have a few others and I'm glad that they're here. So, caucus engagement. When does the caucus meet? Meet once a year?

OZAN SAHIN: Fred, before that, I just have one more reminder. I've just circulated the attendance sheet. If you can circle your name on the attendance sheet, or if you can't find your name, please if you can just add your name anywhere on the sheet, that would be appreciated. Thank you.

FRED BAKER: Okay, cool. So, when does the caucus meet? Well, we meet at one ICANN meeting every year and we meet at every other IETF meeting, this particular meeting being an example.

Matt, did you want to say something from the Membership Committee?

MATT WEINBERG: I can do it now. I think it's 4B on the—

FRED BAKER: Okay. Well, I just did 4A.

MATT WEINBERG: Oh, okay. Sure. Sorry about that. My name is Matt Weinberg. I'm the membership chair for RSSAC. I am joined by Dave Lawrence from

Akamai and Alejandro Acosta who is not here as part of the membership committee.

So, just a little background. The way things work are anybody can apply to the RSSAC Caucus. They submit a statement of interest to the RSSAC and those statements of interest are reviewed by the Membership Committee. Sometimes, they are incomplete or we have further questions, we'll reach back out to the applicant.

If everything looks good and we recommend that they be accepted, we bring it to the RSSAC and the RSSAC formally approves it and then they're part of the caucus.

So, when I joined the Membership Committee, one question I had was we have a lot of people on the caucus. What's the participation level from caucus members?

With the help of ICANN staff, we reviewed participation and participation is multiple things. It's showing up to caucus meetings. It's participating in various work groups associated with RSSAC, publishing a bunch of things. So, we canvassed the data for that and what we found is that some RSSAC Caucus members have contributed and some have never even attended a meeting or published anything or actually done anything with RSSAC whatsoever.

So, what we ended up doing was we're slowly starting to reach out to RSSAC Caucus members who have never participated at all and we did that by reaching out to about 20 people in the past few months.

What we found is sometimes our data was inaccurate. The hadn't, in fact, shown up to meetings and we've since corrected that. Some people responded and said that they are no longer interested in being in the caucus, and for those people, we will remove them. And some people said, "Well, you're right. I haven't done anything. But I will. I promise." I think we're going to keep an eye on that. The idea here is to have caucus members be able to contribute to RSSAC itself.

So, that's where we're at. Dave, did I miss anything?

DAVE LAWRENCE: I'll just say that one reason—I think the most common reason—that people said that they were no longer interested was that the roles changed. They had previously been a member off an RSO or something and were no longer working in the area and that's the main reason they changed. So, the others fell into the category of they're still affiliated with the DNS and do want to continue in some role but they have to re-prioritize.

UNIDENTIFIED MALE: [off mic].

DAVE LAWRENCE: Gosh, I should know that. Ozan, do you have that? It's like 150-some, maybe, plus or minus. I can't remember, to be honest. I should know that by heart but I don't. 125. But it's a steep cliff of participation. There are some members who participated a lot and some members who

we've never heard from ever since joining the caucus. Any questions? Sorry, my back was to the back of the room. Okay, thank you.

FRED BAKER: Yeah. So, that issue from Brad, in my perspective, we'd like to make sure that the caucus consists of people that are helping accomplish the goals, whatever they are. So, that's fundamentally what we're after in that.

Okay 4C, publication acknowledgements. I believe the documents that have been developed by the RSSAC Caucus have in fact included attribution but we're kind of making sure that's happening, and talking about observers and contributors, people that were on the list but didn't say anything, people that did. So, do you want to add to that?

Okay. So, Andrew. Where's Andrew? I don't see Andrew. Sorry, are you virtual, Andrew?

ANDREW: Yes.

FRED BAKER: Okay. So, we have three documents that we're working on right now and Andrew has gone through and made some preliminary comments on but further discussion happening and they're listed on the agenda there. So, let me turn this over to Ozan.

OZAN SAHIN: Ken is talking about [inaudible].

KEN RENARD: Hi, this is Ken Renard, Army Research Lab. Just going to give a quick update on RSSAC 002, some of the efforts. This is not a true work party. It's just a caucus effort. It's going to be done on the mail list. So, in the agenda there, there was two links—one for the statement of work, what we're going to do to update this document and then the other one was the actual document itself.

So, we're soliciting input from the caucus. You can participate by emailing to the caucus list or going into the link that's in the agenda. There's actually a Google Document shared that you can comment on. I'm going to keep the scope of this pretty small, that we're just updating terminology. There's a few things in the statement of work, like measuring QNAME Minimization. That's more of a research reason for doing that. We're going to leave that stuff out.

The things here that were changed, just some of the [inaudible] representation, removing individual IPv6 sources of queries and just doing basically /64s and higher aggregations.

So, the updates. I think there was text that said we will update this every one year, two years. [inaudible] as necessary in order to keep unnecessary work from happening. But please feel free to contribute on the mail list or directly on the Google Doc to make suggestions. Thank you.

FRED BAKER: In general, I would suggest doing it in the Google Doc, otherwise we might lose it. Ozan?

OZAN SAHIN: Thanks, Fred. Regarding RSSAC 023, which is the history document, and RSSAC 026 which is Lexicon, we had work sessions at ICANN 66 which was two weeks ago and the next steps on updating those two documents are, for staff—this is Andrew and Danielle—to incorporate the comments and product the next version of the document and they'll circulate on the RSSAC Caucus list, so that further comments can be received and calls will be scheduled to update those two documents.

FRED BAKER: Okay. So, [inaudible].

OZAN SAHIN: Yeah, 23 and 26. Yes.

FRED BAKER: Do you want to do that?

OZAN SAHIN: So, regarding 23 and 26, these were the two updates on the next steps but if you have any particular questions, you can share that with us or on the mailing list, so please let us know.

FRED BAKER: Okay. So, the issue there, 23 is a history document, and like every history document, it's becoming historical. And we're asking the RSOs to go and update it with respect to things that they have going on.

In the caucus, please feel free to comment there. The principle thing that we're talking about is if the document really stops in 2008 with things that were going on at that time, and time has passed, so let's capture whatever history there is.

PAUL HOFFMAN: Fred, one note is that there is also … Staff has suggested putting in some additional sections that did not make it into 023, so we've already discussed some of that and they may or may not go in. But it's not just bringing up to date, it's adding in other history that didn't make it into 023 as a possibility.

FRED BAKER: Okay. So, you want to inform the room about that?

PAUL HOFFMAN: Sure. The two places that are being discussed, one is to talk about the history of I don't even know the neutral name of when John Postel sort of did a test with the root server operators. The other is escaping my head at the moment. I'm sorry? Thank you. The IANA transition and how that fits into the history.

So, the first one that I mentioned, staff already had written up some. We discussed it at the last meeting. A lot of that is going to get ripped

out and rewritten, so we don't need to discuss that now. Then they're going to write a new section on the IANA transition as it relates to the history of the root server system.

FRED BAKER: Okay, great. So, you will see comments on the mailing list providing a link and saying, "Please go read it," and when you see that, please go read it. Ozan, you look like you're ready to say something.

OZAN SAHIN: In the last RSSAC meeting, I also recall the co-chair, Brad Verd, invited all RSO representatives to reach out to staff and share their individual portion of history, if they have any, with the staff, while updating the history document.

BRAD VERD: 26 is the Lexicon

FRED BAKER: I'm sorry?

BRAD VERD: 26 is the Lexicon.

FRED BAKER: Yeah. I'm going to start touching Ozan's computer here.

UNIDENTIFIED MALE: Fred, I can do the Lexicon one as well. So, the basic changes on the Lexicon were to add terms that we came up with during the metrics work. So that won't finish until we're finished with the metrics work. And then also to copy the exact words of what is RSSAC and what is Root-Ops from the recent RSSAC document on those. So, that would just mean if someone sees a new Lexicon, they'll have a slightly larger view.

The main change from the first one was Anycast and Anycast instances and capitalization and such like that. So that was more of a housekeeping one. But the new terms are things we came up with in metrics and things that RSSAC came up with in defining what is RSSAC and what is Root-Ops.

FRED BAKER: Great. Thank you. So, any comments on that? Brad?

BRAD VERD: Really quick, I guess, for edification of the room, there's only a handful of documents that we've kind of versioned up. Our normal rule of thumb is to do a new number for a new document. But a number of these documents are kind of living documents, the history document, the Lexicon. We're not going to make a new number each time we add something to the Lexicon. So, I just want to make sure everybody is aware of the normal practice.

FRED BAKER: Okay. So, if there are no further comments on that, I want to move on to the current work parties and work products in production. Duane, do you want to talk about the metrics work? I'm sorry, I skipped something. Believe it or not, we had a workshop, and at the workshop, for the most part, what we talked about was metrics. But we also had a few other topics in there.

So, when ICANN spent a whole lot of money to bring us all to somewhere and talk in a room, we kind of feel like they should get something back. So, the something back is a workshop report. So, there is, in fact, a workshop report in—what number is this? Is it 44? Yeah. So, RSSAC 44 is a workshop report from a workshop that we had about a month ago, like I say primarily talking about metrics but also a couple of other things as long as we were in the room. That's to bring people up to speed.

Okay. So, now, having done six, seven. Work parties and work products. We have two work parties going on, Metrics and the Resolver Study Work Party. Duane, do you want to step into that?

DUANE WESSELS: So, while Ozan is getting the slides up, I'm Duane. Russ Mundy and I are cochairing or co-something-ing this work party. It feels like it's been going on forever, but I think it's maybe only about a year.

So I have a presentation. This is really kind of a high level presentation of the work party and of the document. I hadn't really planned on getting into some of the nitty-gritty details about what's in the document. I would encourage people to go and look at that and either

make comments in the document itself or on the caucus lists. I guess if we have time, we can maybe address some specific things, but for the most part, I just wanted to run through this.

So this slide here shows the high level structure of the metrics document. There are these nine sections and I'll go through those in more detail, one by one. So next slide.

The introduction contains a copy of the statement of work to kick this whole work party off, and then it also has another little section that just outlines the whole document structure and it's got your standard RSSAC document boiler plate, which says this is a report of the work party and so on, and so on, and so on.

So the most interesting thing there is probably the statement of work. There has recently been some new text proposed in the section so definitely take a look at Section 1. If you're paying attention to this document, definitely take a look at the introduction and make sure that you're up to date with what's there and if you have any comments, please leave comments in that section. Next, please.

Section 2 is called Background and Scope, and I spent some time talking about what sort of things are in scope versus out of scope for this work. References that the purpose of this document is to talk about minimum levels of performance for root servers and the root server system. Early on, a long time ago, when we were starting this, we weren't as clear on this point and we also talked about maybe doing what we call good levels of performance but that has been taken out and now we're focused only on minimum levels of performance.

This section references the RSSAC 037 governance document and makes reference to some of the things that are called out in that document, so the performance, the PMMF is the Performance Monitoring and Management Function or something. And this section also makes reference to the phrase, SLEs, Service Level Expectations.

In Section 2, we talk about things that are not in scope such as these metrics are not being designed for conducting research and performance trends, not designed for making comparisons between RSIs, which the RSI here is Root Server Identity. This is kind of a new phrase that we're throwing into this document. We don't talk about RSOs as much as we talk about RSIs which is the identity, or formerly, letters of the root server operators or the root servers.

There's a section on prior work which references a number of RSSAC documents such as 001, 002, some other numbers, but various RSSAC documents and an RFC or two. And as Paul was just mentioning, there's a terminology document which now, we're in sync, we've got this terminology and the RSSAC, the new lexicon terms in sync. There are some new terms to find here and I guess those will also appear in the lexicon as necessary. Next, please.

So yeah, the terminology section has these definitions. Measurement means a single, sort of a smallest unit of measurement. It's a single set of query response. Metric is an aggregation of a number of measurements together over some period of time. Threshold is a value of which the metrics are applied. If you exceed the threshold, it's bad. If you're below the threshold, it's good, generally. There's a definition for a vantage point. Formerly, these were also called probes but now we're

very consistent in calling them vantage points. This is either like a software or hardware device that does measurements and we also have a definition for collection system which is the thing that receives all of the measurements and aggregates them together and prepares reports and things like that. Next.

As far as vantage points goes, this document makes a number of recommendations on those. It says that there should be approximately 20 and that they should be distributed approximately evenly among five geographic regions. In this section also, we note that this is just sort of a first attempt and maybe not the best way of doing things and so there's a future work item maybe to do better here and to come up with better ways of figuring out where vantage points need to be located and that sort of thing.

There is some text in there that talks about how vantage points are connected, that they should be placed at well-connected data centers with good power and so on, and that vantage points in the same region should use different connectivity providers to avoid certain measurement biases.

There's this long section with these ten or so sub-sections called general aspects of measurements and metrics and I won't spend a lot of time on all of these, but in particular, I would call your attention to some of the latter ones. There's, I think this is 4.9, I believe. So 4.9 is called "Determining the Number of RSIs required for reliable operation of the RSS". This has some mathematical formulas in it and it's kind of the basis for a lot of the other things that appear later in the document, so

this is a very important piece to understand and make sure that we have consensus on.

There is also, the last one, I would call your attention to it only because it's sort of new so if you haven't seen the document recently, you maybe want to take a look at potential effects of metrics on independence and diversity. The other sections have more or less been stable for a long time, I believe, and there's probably not a lot more to say about that I guess. Next, please.

So this is that section on determining the number of root server instances. So I think the specific language in there, it says something like, "determining the number of root server instances required for reliable operation of the root server system". And so this is some formulas and discussions that we had at the RSSAC workshop that Fred was just referencing. We came up with this formula and sort of agreed on all of this. What this formula here means is that if you're sending a query and if you're first, you have to send a query to a server that's down, so the first server you try is down, the second query should be successful with a two-thirds probability.

And in order for that to be true, and if we have 13 root server identities, then we should have eight, K equals eight, is the number that comes out of this formula. There should be eight identities up at any given time in order to provide reliable operation of the root service. And so we'll see this K equals eight and N and K throughout the document later in a number of places.

Any questions about this? Because this is kind of an important thing. I want to make sure that everyone sort of gets where we're coming from here because a lot of the metrics and thresholds derive from this. Okay. You can ask later if you want to. Yeah, go to the mic, please.

PAUL MUCHENE: So why two-thirds, not maybe three-quarters? What drove the formula or how are you thinking about it?

DUANE WESSELS: So the number two-thirds was just, in some sense, it's a number that's made up. It's what the people in the room felt was an appropriate value for how we think the root server system should work reliably. There's no other real solid basis for it other than that. It's our expertise and what we agreed on at the time.

Okay. Next?

So this chart actually appears in the document and those of you paying attention will know that it's fixed now. In the previous presentation, this was incorrect. The reason that this chart is included is to highlight the fact that this value of K equals eight depends on the number of N equals 13. So if the number of N changes in the future, then K should be updated appropriately as necessary. And so this shows the values of K for different values of N and the red bar highlights where we are today with N equals 13.

So in the document, Section 5 is describing these four root server identity metrics of availability, response latency, correctness, and publication latency. Next, please.

So root server identity availability is measured separately over all the combinations of v4, v6 and UDP TCP transport. This is a pretty straight-forward measurement. You send some number of Q queries and count some number of R responses and then the availability is just the ratio of those two. It's R divided by Q.

Remember that these are reported over a period of a month, so all those measurements over that month period would be aggregated together and then its availability calculated. The threshold for this metric is 96% and I believe the next slide maybe explains how we got to 96%. Yes.

So the way that we got to this 96% was applying this standard formula of what you could call simple, parallel K out of N availability and that's represented in this kind of gnarly looking equation here. But this is kind of a standard thing that you could see in reliability and availability textbooks and things like that. The way we get to 96% is that we say, "Well, we want the overall availability, the capital letter A, to be five nines, 99.999%." We have N equals 13, and K equals eight. And that tells us that the little a needs to be 96%, 0.96, to meet that requirement.

Now there's a lot of simplifications here because this formula assumes that all of these individual components, in this case, all of the root servers, are identical, that they're independent and that sort of thing, and we know that's not the case here but we're applying this model and

this formula regardless because in some sense, it's the best we can do and it's relatively easy. But there are certain simplifying assumptions here which we need to keep in mind as we look at the data.

UNIDENTIFIED MALE: So question for you. We have a person in the chat room over there. Do you want me to bring them up?

DUANE WESSELS: Sure.

UNIDENTIFIED MALE: Okay, Shailesh Gupta, you are online.

UNIDENTIFIED MALE: He's trying to speak and he's having audio problems.

UNIDENTIFIED MALE: Okay, I think you're speaking but we can barely hear you.

DUANE WESSELS: Can he type his question? Is there a chat function?

UNIDENTIFIED MALE: Okay, we can't hear you.

DUANE WESSELS: Yeah, the audio's not good enough that we can understand anything, so apologies.

UNIDENTIFIED MALE: And I'll drop it.

DUANE WESSELS: Okay.

UNIDENTIFIED MALE: You could ask him to type it into the chat.

DUANE WESSELS: Yeah, I don't know if it has that feature or not. So [inaudible], if you can get the question to us, we'll answer it for sure. Meanwhile, I guess we'll proceed. Can you go to the next slide, Ozan?

Okay, so that was root server identity availability. The next metric is response latency. As the previous one, this is also measured separately for v4, v6, and UDP and TCP. In fact, it uses the same query or the same measurement that was used for the other one so this is nice. This is measurement reuse. We don't have to send a separate measurement for this metric.

Here, of course, what we're interested in is the latency of a query and response. How long did it take and the metric value is calculated as the median value of all of the roundtrip times from all of the probes over the one month period. So given our measurement interval of five

minutes and approximately 20 vantage points, that means in a typical month, you would have about 172,000 individual measurements aggregated together.

And the threshold for this, again, the threshold that the median value is 250 milliseconds for UDP and 500 milliseconds for TCP. The rationale for TCP being twice UDP is that there's this connection set-up latency to account for.

PAUL HOFFMAN: Just to be clear since other people may not be seeing it from the slide, it's measured separately for v4, v6, UDP, TCP, and the thresholds are also separate for them.

DUANE WESSELS: That's right. It's measured separately. It's reported separately. But all the thresholds, the UPD thresholds are 250 and the TCP thresholds are 500. Yeah.

Okay, next.

So this is some data from RIPE ATLAS that I was looking at last month and included here. This is from a single, unnamed root server identity, but it shows the distribution of latency measurements from RIPE ATLAS anchors and so in this case, it's not drawn in here, but if you imagine a horizontal line at the 50th percentile, where it intersects those green and purple IPv4 and IPv6 distributions, it's somewhere in the range of 50, 60, 70, 80 milliseconds, something like that.

The vertical bar shows the 250 millisecond threshold for UPD. I don't think it says but this is probably for UDP only. This is not TCP because mostly what RIPE ATLAS measures is UDP.

All right, let's go to the next.

The next metric in the document is correctness, so root server identity correctness. This one is a little bit more complicated than some of the others. There's a lot of stuff to take in when you read this particular description of the metric. It's based on both what we call matching and validation. So matching means looking at the R data values and that they match a value published in a root zone file, for example, whereas validation means you do sort of standard DNSSEC validation on the response data that you get.

Correctness is not measured separately per transport. When a measurement is done in this one, the transport is sort of chosen at random and so instead of there being sort of four thresholds or four metrics reported, in this one, there's just a single value. This metric says that 90% of the measurements should be done for what's called existing data, so data that's in the root zone. This is your DS records, your NS records, your DNSQ records and so on, and 10% should be for non-existent data. And it specifies the format of a query name that would be used to generate a query expected to be non-existent. You see in here it says, RSSAC 0XX. That will be the name and number of this document when it's published and $ random is a ten-character random string, so ten random characters. That's designed to sort of cover the whole name space of the root zone and check for things like proper [NSEC] responses and whatnot.

The calculation of the aggregated data is just the number of correct responses divided by the total number of responses. And the threshold for this metric is 100%, so every measurement made by the system should come back as correct and if it's not, then that's a problem to be investigated later and identified as a potential problem.

I think, as I said, it's kind of a complicated thing so I think Paul in particular would like a set of eyes on the rules. There's a set of rules for how you go through and validate the different types of responses you can get. Most of this was written by Paul and I think he would appreciate additional people to look at it and make sure that we got it right because it is nontrivial. All right, next please.

The last metric for an individual root server identity is the publication latency. The idea here is to calculate the time taken between when a root zone is made available from the root zone maintainer until it's picked up by the root server instances and served to the users. So this one also reuses the queries from the availability metric. These are SOA queries and we look at the serial numbers in those responses. Since those are done on a five-minute interval, that's our maximum resolution is five minutes.

There's an important role for the collection system to play so that the central collection system has to sort of know when new root zones get published by the root zone maintainer, and that can be done by looking at the serial numbers in aggregate and saying, "Oh yeah, I see this new serial number starting at this time so it was probably published close to that time." And once it has that time, then it can calculate the amount of time taken for the individual identities to start serving the new zones.

In the metric, the latency, here it says median. Paul, I believe you had a proposal to change this to mean maybe. But so far, we've considered this as a median value.

PAUL HOFFMAN: It was actually Ken who, someone else pointed it out but I agreed. Ken, didn't you point it out and then delete?

KEN RENARD: Sure.

PAUL HOFFMAN: Okay. So anyways, yeah. So what was pointed out is that if an RSI was just wildly horrible for 48% of the time and then fine for the other 52% of the time, if it was a median, then that would be considered okay. If it's a mean, then that would definitely knock them down which is what I believe the work party would want to do in the case of where almost half of their responses or almost half of their latency was two days.

DUANE WESSELS: Okay. I think we'll have to think about that.

UNIDENTIFIED MALE: Now question for you. So a mean depends on it having a [gousean] distribution in the random number sequence. Are we saying that we think that this is [gousean]?

PAUL HOFFMAN: It doesn't depend on it being [gousean] for a mean. I mean, it's just a straight …

UNIDENTIFIED MALE: Well, yeah. You can take all the numbers and average them. But as far as the mathematics, what you're looking at is a [gousean] distribution.

PAUL HOFFMAN: I'm sorry. Can you say it again?

UNIDENTIFIED MALE: What you're looking at is a [gousean] distribution. It's the center of the [inaudible].

PAUL HOFFMAN: So at least in the initial attempt that Duane and I looked at, it was not [gousean] at all. It was very, very good except for occasionally where it was slightly bad. Partially, that's due to the five minute granularity, but the other is due to it depends on which instance of an RSI happens to be closest to the vantage point. And given that each root server operator updates their instances differently, some of them will push them all out instantly, some are waiting and such like that, it really wouldn't be the kind of [gousean] curve that you would think. So that's why a mean of just a plain, old mathematical mean seems reasonable, or at least, it seems more reasonable than a median which could gloss over large

problems that happen enough that we would be concerned. It may not be. It may that the work party actually is fine with having an RSI be wildly off for almost half the time but not quite, and then that still would be considered to be meeting the threshold.

[DALE]: So my main thought about this is I don't want to diminish the significance of the argument about mean versus median. So I'm not trying to trivialize it with my next suggestion which is it's also considered that we look at what's the eventual use of having these numbers and these statistics. If it's important on an operational basis, it seems to me that this type of really exceptional circumstance where an RSI is great most of the time except for the other almost most of the time where it's terrible, that we're going to have many other means of noticing that something is really erratic going on. And so I don't think we're relying on this data to achieve that, and again, I'm not trying to minimize. I do think it's a meaningful discussion to be having but I'd like to keep in mind the end goal of just how important it is to settle it in some ideally perfect way versus a pragmatically useful way.

DUANE WESSELS: All right. Okay, thank you. Yeah, so we'll leave that for future discussion on this particular question. I guess one thing that bugs me a little bit I guess is this is the only, if most of the metrics are medians and this is the only one that's a mean, so it's a little bit strange that way. But I could live with that.

Oh, and so lastly, the threshold for this is, I forget. We sort of went back and forth between 60-65 minutes. The idea is that 60 minutes is twice the SOA refresh value and then there's a five minute wiggle room because that's sort of the granularity of the measurements. Did we leave it at 65 for now? Okay, so that's what it actually says. Ray?

RAY BELLIS: Something's occurred to me the last couple of days and I put a note in the doc so we can maybe talk about this one on the next call. And it's more actually, I guess, on the correctness rather than the publication latency but this kind of relates anyway.

So when buying starts up, if it's got a copy of its own file cached, it will load that immediately and then it'll go through a refresh cycle. So you've got a window when buying starts up, when you might be serving stale data. As I understand it, maybe Andre can confirm this, if it detects that the zone file is older than the zone expiry time, it won't actually load that zone straight away, so we've still got this mismatch between the zone age and the refresh and retry because under just normal operational circumstances, you start with a server that's been down for four days, it will come up with four-day old data. And I don't think that's been accounted for currently in our view of what's correct in zone data.

DUANE WESSELS: So yeah, I think you're right because the correctness metric says that you go back in time up to two days and the data that you get back should be no more than two days old. So if that's something that happened, then you would have a case where –

RAY BELLIS: It's a short window. It's a [inaudible] window, but it can happen.

DUANE WESSELS: It can happen, yeah.

PAUL HOFFMAN: I have a question on that. So we don't know when it would be [inaudible].

UNIDENTIFIED MALE: Speak into the mic, please.

PAUL HOFFMAN: I'm sorry. I was. We don't know when, in that case where it has started up using its cache, it won't ever look again until it gets pushed or somehow notified by, in this case, the RSI because we're talking about instances here.

DUANE WESSELS: Yeah.

RAY BELLIS: Actually, I think it will automatically start a [inaudible] refresh anyway at startup but I don't think it will do that until the zone data has already been loaded and made available.

DUANE WESSELS: Thanks, Ray. So I think this is probably a good time to point out that throughout all this metric stuff, we're definitely giving ourselves lots of chances to make mistakes and fix them later, right? So if this is something that we find is a problem then we can, if we don't fix it now, then we can definitely fix it later. We can address it in a future version of the document if we need to.

Okay, move on to the next section.

So Section 6 in the document is about RSS metrics, root server system metrics, and it has the same four metrics, the same four sub-sections, availability, response latency correctness, and publication latency.

RSS availability is defined quite a bit differently for the RSS than it is for an individual RSI. This is a case where we use our value of K, K equals eight, to measure this. So the ways this is currently described, for every time interval, T, and from every vantage point, V, you would determine R_{TV} which is the number of root server identities that respond to an SOA query.

Now fi that value of R happens to be greater than K, greater than eight, then you would clamp it at eight, and in this formula, that's what the numerator says. You take the minimum of K and R and then you sum up all those, and then you divide it by a sum of K for every measurement that you have. So that's essentially K times the number of intervals times the number of vantage points. And that's your root server availability. Or that's your RSS availability. Excuse me.

UNIDENTIFIED MALE: Four of them.

DUANE WESSELS: Four of them. Right, so you do that for each transport separately. That's right. Again, this reuses the measurements from the other metrics.

And the threshold for this is 99.999%. The reason for that is because that was sort of a requirement for how we got to K equals eight earlier in the document. So I will say that this is something that I've been looking at just recently in some of the data from RIPE ATLAS. I did some analysis to see if we took that RIPE ATLAS data and applied these measurements, do we get the right thing out? And this is one where I'm not sure that we do. The data didn't quite come out the way I expected, so I think this requires a little more investigation and I want to get a second opinion. One of the questions here is did we not get what we expected because of the way RIPE ATLAS works or because of the way the measurement is defined, the metric is defined? Or is it something about the way the root servers are being operated? So there's these three unknowns and if we repeat that experiment with other data besides RIPE ATLAS, we may get something that makes more sense.

Paul has done a proof concept implementation on these metrics. You might have a month by now. You'll be getting close maybe. But as Paul collects more data, we can use that and there may be other places we can get data to look at this. Go to the next slide.

I think there's a table. Yeah, so this, I forgot to fix this. I apologize. This is hard to read because this is basically an image cut and paste from the document. But this table appears in the document and it shows some examples of the RSS availability you would get in certain really contrived situations.

So for example, if there happened to be an attack or some other reason that for a whole month, one root server identity was out, was unavailable, if it's just one then you still get 100% availability because in every interval, you would measure R equals 12, and that's well above eight, so you're still at 100%. If there's another attack or whatever reason that five root server identities are unavailable for the whole month, you still get 100% availability based on this formula.

However, if you lose six identities entirely for that whole month period, then the measured availability is 87.5% which is seven-eighths because in all those intervals, you measured seven and you needed eight in order to meet your 100% requirement.

If there was a 24-hour tack that took out all of the identities, then the availability would be 96.6% which is 29 divided by 30, so 29 days out of 30 days is how you get that. If there was one five-minute interval in which one vantage point can only reach seven root servers, but in all the other intervals, you could reach eight or more, then you have 99.9999%, so six nines availability. If in two intervals, you had only seven vantage points, then you get close to, I'm sorry, two intervals, only seven vantage points can reach no root servers, then you're basically at the five nines availability level.

Any questions about these contrived examples or complicated mathematics? Okay.

RSS response latency, here again, we're using the K from our previous formulas. So in each interval, and from each vantage point, you find the K lowest latencies and consider those as a subset of the latencies and then you calculate the median of that subset. So since this is a subset, we expect that the median value should be lower than for the RSI case, so the thresholds here, 150 milliseconds for UDP and 300 milliseconds for TCP.

[DALE]: So on the previous side, I actually did have some thoughts related to it that I think are maybe not worth extended discussion right now but could lead to some interesting discussions, which in particular, is an attack that took out six root servers for a month is a pretty remarkable attack, right? I have to wonder what kind of news attention this would be getting, what kind of Internet performance, and I know the purpose of this work is not to figure out what's the end user implications of this, but I don't have the good intuitive sense of whether such a remarkable attack but then still saying, "Well, we've had seven availability is a good number or a bad number for that," right? On the one hand, I think it's tremendous because it shows just how resilient the system is, that with not exactly half, but half of the instances up, we're still in really good shape. But on the other hand, six taken out for the month is just so remarkable that saying I'm not really sure what sense to make out of the results of these contrived scenarios.

DUANE WESSELS: Right, and that's why I tried to make the point that they are contrived. And it doesn't have to be an attack either. It could be something else, going out of business or whatever, some other reason, fiber cut –

UNIDENTIFIED MALE: [off mic]

DUANE WESSELS: Yeah, it could be other reasons that they're out. But I definitely did want to make that point that based on our definitions and our formulas here, this is what it's telling you that if you had a situation where the root server system was down to that capacity, you still get 100% out of these, or you get whatever, 87% out of these formulas.

RUSS MUNDY: One of the things that we might or might not want to include when we actually publish it is this table or something like it. I think this is a very important table to have in there now so we can see what the mathematics translates to, but it might be something we want to discuss before actual final publication, if we ought to have something like this in the document itself or not.

DUANE WESSELS: Or reword it if necessary.

RUSS MUNDY: Yeah.

DUANE WESSELS: To not talk about attacks that take out things for a whole month. Yeah, okay. Thanks.

Okay. All right, so back to RSS response latency, so again, it's the K lowest latencies in each interval aggregated together. The thresholds are 150 and 300 milliseconds, pretty straight-forward.

The RSS correctness metric is a simple aggregation of all the measurements from all the individual root servers and the threshold, again, is 100% so the root server system needs to be as correct as the individual identities need to be. There's no tolerance there for mistakes. And next.

The RSS publication latency is also a straight-forward aggregation of the measurements from all the individual root servers. Here, again to calculate median, maybe mean if that's what we settle on, some average of all the publication latencies from all the vantage points for all the root servers. You'll notice in here the threshold is lower and the consensus was that it would be very remarkable or very bad if it was the situation that so many root servers were having such a high publication latency at the same time that you needed the 65 or 60 minute threshold. So the expectation here is that a majority of the root servers will be having very good publication latency and so this threshold can be lower than the case for an individual root server identity.

All right. So I want to make a few points about how the results are reported. I think this actually appears earlier in the document, but anyway, for the RSI metrics, for the root server identity metrics, the reports that come out of this work are designed to just be pass/fail, either the threshold was met or it was not met. However, for the RSS metrics, the actual numbers would be reported, the actual latency numbers, the actual correctness numbers, and so on.

Furthermore, in the interest of transparency, the raw data is to be made available to anyone who seeks it out or asks for it. So even though the reports have just pass/fail, anyone willing to put in the time and the extra effort and dig through the data, they can do that and do that on their own.

Okay. I have some slides on the recommendations, and the reason I put this in here is because these were changing a little bit towards the end of our work last week and before that. So I feel like these are important to get before people and make sure that we have consensus on.

I'm going to read this, but before I do that, the last time we talked about this, we talked about the word "initial". The word "initial" appeared in the previous definition but I think it said "initial operational experience" rather than "initial implementation" so just to point out to people that we are aware that the word "initial" appears here. When Russ and I sat down with the staff and worked on this, we really couldn't find a way to say what we wanted to say without using a word like "initial" so that's why it's here.

Okay, so here's the recommendation. The RSSAC recommends that the ICANN Board Commission and initial implementation of the measurement system described in the stock [event] to gather operational data and experience from actual monitoring of the RSS. The initial implementation should be designed such that it can transform into the official implementation as described in Recommendation 2 below. The insights learned will inform future revisions of this document if necessary.

[DALE]: I think one of the problems with the original wording, proof of concept, right? I think "initial" is not so much a problem. What's the controversy over "initial" now?

DUANE WESSELS: I don't remember, to be honest. But proof of concept was there before. You're right. We took that out. Somebody had some issue with "initial" and so we took it out before, but we felt like we had to put it back in.

[DALE]: Okay. For the record, I'm okay with "initial". "Proof of concept" bugged me but "initial" is good.

DUANE WESSELS: All right, thanks. And so if you happen to go through the Google Doc and see this, you'll see this Recommendation 1 labeled as something like New Recommendation, Proposed New Recommendation 1, and the old

one is above it. So they're both there so you can compare but this is the proposed new one that we hope goes forward. Okay, next.

Recommendation 2 is sort of this bullet list. To be clear, these sub-bullets here are abbreviated from the Google Doc so if you're a stickler for details, make sure you read the doc. These are just sort of to capture the essence of it. But this talks about the official implementation of the metrics and has these requirements. So it meets the requirements for the vantage points as specified. The software is published as open source, the raw measurements. The raw data is made available to anyone interested in transparency. Monthly reports are published pass/fail. The methods and whatnot must be described and that in the case where thresholds result in some kind of failure or whatever, those underlying results have to be shared with the root server operator to investigate and understand what's really going on there.

I think this has been relatively stable over time. There may be just a couple of little minor wording tweaks here but again, since these are the recommendations, I wanted to get people to look at this and make sure that we all like it. Okay, next.

And Recommendation 3 talks about possible future work. So it says the RSSAC, in collaboration with ICANN in the Internet community, should consider the following additional work. This is not a "will consider as future work". This is "should" or "may consider as future work". So one of the things, the first bullet talks about something that we talked about a long time ago in the work party, which is the trade-offs between doing third party measurements, so measurements that can be done remotely by a third party versus measurements that need to be self-reported and

done by the operator themselves. And so this references that if we wanted to get a more holistic view of the system, it may be necessary to have the root server operators do their own measurements against all of their instances and then self-report those. But that's not considered in this first version.

Recommendation 2, make a reference data set, something that we still plan to work on. The idea here is that if you have a reference data set, then you can write your own code to process that and make sure that we all get the same result when we apply the formulas and the methods and the metrics and make sure that we're all on the same page there.

The third bullet says, it talks about exploring the financial aspects of increased accountability and how it relates to these metrics. We have a little bit of that already in the introduction but that was something that some people felt was very important. Similarly, to publish a document that advises any bodies created as a part of the ongoing governance work on how they should interpret and use the data from the measurement system.

The next one talks about the better long-term plan for locations of the vantage points which we mentioned before. And lastly, some work party members felt that future versions should maybe take a more analytical approach to modeling whereas what we have now is sort of more empirical. So that's another possible future work.

Next. So Section, this might actually be Section 9 now but there's a section of example results and the idea here is just to show one way that the reports might be presented to match the expectations from the

work party document. So this one shows the metric of RSI availability for hypothetical end route over the different transports of IPv4 IDP and so on. The green and red show the cases where the thresholds were met in green and not met in red. And the last column shows the number of measurements used in the aggregation. This is something we didn't really talk about today, but in the document, it says that in order to convey the precision of the measurements, you have to say how many measurements were included in the aggregated interval.

Yeah, that's it. I don't know how much. How much time do we have for Q&A versus the rest? Who wants to talk metrics some more? Come on.

FRED BAKER: Well, we're very close to the end of the agenda, so if you've got questions, go for it. If there's one question per minute –

UNIDENTIFIED MALE: [off mic]

DUANE WESSELS: Okay. All right. I would just, again, encourage people to, if you have comments or questions, go to the Google Doc. If you need the link, find one of us or find the staff and get the link to the document. Read it through. We're trying to wrap this up as hard as we can. We hope to get it finalized and get a final version out to the caucus. And then after that happens, there would be a version that goes to the RSSAC for voting and approval, maybe in the next couple of months if we're lucky.

UNIDENTIFIED MALE: [off mic]

DUANE WESSELS: Yeah, so Russ said to mention another call. Given the fact that we did have some recent additions and some text where there's still some debate going on and some of the other things, Russ and I think that it would be good to have another metrics work party call. I know it's kind of not the best time of year for that, but we'll figure something out for scheduling the next call and the goal with that call will be to resolve the outstanding issues in the document, go through them one by one. So look for that. All right. Okay.

FRED BAKER: Thanks, Duane. So the one real thing left on the agenda is talking about the resolver study work party, which I'm the shepherd for and Pauls' the one that's doing the actual work. So Paul, do you want to comment on that?

PAUL HOFFMAN: Sure. As we talked about in the last meeting, pretty much work is done on that and there has not been a whole lot of people using it. However, today Paul Muchene and I have been working on it, found a major error, already fixed and already in the repo. So the next step for this would be for me to, unless there's a lot of interest all of a sudden that comes up in the next week or so, which I'm not expecting, for me to write up what was the statement of work, what came out of it. The repos are open. At

the last meeting, I mentioned that I had an issue where I realized I had put a license in on these, a very open BSD three clause license and had not gotten that covered with ICANN's legal department.

They looked at it and said, "Oh yeah. No, we always do that." So that was good. So basically, that'll get done in the next few weeks and then we can call that done for now. It's definitely ongoing work if people want to do it. The repos are open, poll requests are open such like that. As people start using it to do experiments on their own to figure out, to actually use the test bed to get interesting information, it may be that there are extensions to the test bed that people want. We don't need to open another work party or whatever. We can just be extending it. My hope is if we get some interesting questions that have answers that come out of the test bed, they come to the caucus and maybe collect them at some point in a document. But I don't want to get ahead because we actually aren't there at all. But basically, the work is done and we'll have a closing report for it.

FRED BAKER: Okay, thanks Paul. Anybody have questions for Paul? Comments?

RUSS MUNDY: This is something I hadn't really even thought about until I was just hearing Paul's report but we, in the caucus, have a tools repository. Is this in some way going to have pointers to or connections with the tools repository that's already set up or will it remain sort of a separate activity?

PAUL HOFFMAN: The answer is it could be both. That is it should probably remain in the repos that it is now. I don't know, and Wes isn't here, I don't know if the tools repository actually is meant to be active or to be a collection of tools. I'm hoping this stays active. That is I really do think that since I haven't had enough people using it, I don't know what the next person is going to want from it and it's easily extensible so I want this to be sort of more live. But we could certainly put something in the tools repository saying, "And by the way, this work party did this and here are the links." And basically, the report that we do could be there as a pointer for it.

RUSS MUNDY: Thanks.

FRED BAKER: Okay. Anybody else? So let me call for any other business. Seeing a stampede at the mics, I think I'm going to adjourn the meeting. Thank you.

[END OF TRANSCRIPTION]