SINGAPORE – RSSAC Update to the ICANN Community and Key Discussion Topics
Monday, March 24th 2014 – 16:30 to 17:15
ICANN – Singapore, Singapore

UNIDENTIFIED MALE: March 24, 4:30 p.m., RSSAC update to the ICANN Community and Key Discussion.

LARS-JOHAN LIMAN: Alright, ladies and gentlemen. I would like to welcome you to the session on the RSSAC overview and reorganization process. This is a public update from the Root Server System Advisory Committee.

My name is Lars-Johan Liman. I'm one of the two co-Chairs of the Root Server System Advisory Committee. I share this burden with Jun Murai from the WIDE project in Japan. He was, unfortunately, not able to make it to this meeting, so I will be introducing our [speakers here] and also speak a bit myself about this.

As you may know, the RSSAC is undergoing a reorganization process. But first, I would like to ask Suzanne to give an overview of what the RSSAC is and what our place in ICANN is. So please, Suzanne Woolf.

SUZANNE WOOLF: Sure. Thanks, Liman. We wanted to also introduce our colleagues who are here.

LARS-JOHAN LIMAN: Yes, we should do that. RSSAC is, in its current state, primarily the executive committee. So I would like to introduce the executive committee. We consist of representatives from all the root server operators and a number of liaisons from other groups.

So with me here today: presenting is Suzanne Woolf from F-root and Bill Manning from B-root. And I would also like the rest of the executive committee who is here in the room – those of you who are here – to stand up so I can introduce you.

It's Marc Blanchet who is the liaison from the IAB. Elise is the liaison from the IANA. Duane Wessels, liaison from Verisign. Brad Verd, representative from A- and J-root. John Crain, L-root. [Tripti Singhal] from D-root. Who else? Russ Mundy is the liaison to the SSAC. And who else? [Hirohito] is representative from M-root in Japan.

So those of us are here, and please approach us afterwards if you have questions and want to talk to us. We're happy to talk to you. So, Suzanne, please.

SUZANNE WOOLF: Thanks, Liman. I'm looking around, and we have some smiling faces. Not too many, and you don't hear this very much at an ICANN meeting. We actually like it when we're low-profile and people don't seem too worried about us.

I'm going to do a real quick overview of who we are and what we do, and then we're going to talk a little bit about the reorganization we're going through and also some of the technical work that we're doing with the root server operators in the community.

I was asked to start with who we are in the ecosystem, because as liaison to the Board and one of the few RSSAC people who has been regularly attending ICANN meetings for a long time, I've done this rap.

Just real briefly, though, who we are: the initial membership and key membership of RSSAC includes the root server operators. All of the organizations that operate root servers are represented on the RSSAC exec.

In addition, Liman will talk a little bit more about the reorganization and how we're doing this more broadly, but we also work with other stakeholders from the community with a direct interest in a stable, robust root zone and root data distribution system.

So we maintain a technical and operational orientation. We are not a policy body. We look primarily to DNS experts, to the ccTLD and gTLD (it should have said also) and our technical folks who are our principle customers and partners. We are working towards broader multi-stakeholder representation and interaction. So that's actually some open topics there, how we work with other groups in the ecosystem.

Apologies for the org chart, but we don't even always know, so we thought we should include the graphical version of where we sit in the ICANN ecosystem. We are an advisory committee, like several others: like the ALAC, like the SSAC.

What we're here to do as operators, the root server folks are here to make sure that DNS queries to the root are answered as quickly and efficiently as possible with the most up-to-date information available is distributed by the root zone management partners.

We are here to tell people what we do and to work with the community to do outreach. We provide liaisons and support to, particularly, the NomCom and the Board, but other roles and task forces as requested.

And we're here to provide DNS expertise and advice on a slightly broader basis to the Board and the staff and the community on issues such as DNSSEC and how it should be integrated into the root; root scaling as part of the new gTLD inquiries and preparation; and similar technically- and operationally-oriented DNS issues.

We also work with the root server operators and with the broader community and we will be talking about some of the substantive work underway on measurements and service expectations as relates to the root name servers.

We sometimes find it useful to also point out to people what we don't do. We are not the place to go to tell root server operators what to do. We hope and expect that a role in forming the advice means that we're forming advice that root server operators and other partners can actually take and operationalize. But it's a collaborative effort and a collaborative set of relationships.

We don't tell ICANN what to do. Some of the ACs and SOs have a policy function where the results are binding on ICANN, on the Board, and the organization. That's also not part of our role.

We strictly support and respect the difference between the policy relating to what goes into the root zone and the operations of delivering the data to the users of the Internet. So we set no policy; we have no policy role regarding the contents of the root zone.

In addition, because of the nature of the service we provide and because it must be of the highest quality and provided to all Internet users on an equal and the best possible basis, all of the organizations that operate root servers do other things as part of their business and operations and community activities, but we don't mix root server operations with our other operations. The separation there for operational and management reasons is quite strict. So when you're talking to any of us in our capacity as RSSAC or as root server operators, that's strictly what you're getting.

Further information, just on the technical and operational aspects of root servers: people like that chart because it shows some of the breadth and variety of distribution of root server resources worldwide. There's lots more information at rootservers.org, which is a public site that the operators maintain for public data.

And because it's such an important topic in this meeting, we thought we would say a couple of words on where we see the IANA functions transition process that has been so prominent in conversation leading up to this meeting and already here.

A couple of very simple points for us: most of the root server operators have been acting in that capacity at least as long as ICANN has been around and the current institutional arrangements around IANA have been in place.

And the transition has been – as many people have already said here and will continue to say – has been in the long-term plan from the beginning. We expected it and welcome it and are happy to work on it with the rest of the community.

We do expect to participate in the process along with everyone else, with all the other stakeholders involved in making sure that the future of IANA and the IANA functions is stable and trustworthy for the entire community.

And for us, as we participate in the discussions, consistent with the role that we have always played, the important principle, the thing that we are always here to protect: the security, stability, and resiliency of the root zone distribution system for all Internet users.

Straightforward stuff, real simple, and I can turn it over to Liman to talk a little bit about how we're doing our reorganization to support some of these things within the evolving community.

LARS-JOHAN LIMAN: Thank you, Suzanne. So I will continue a bit with the reorganization. RSSAC, being one of the advisory committees, undergoes periodic reviews as all the bodies of ICANN. And we received feedback from the last one, which we're now trying to act upon and implement.

So we are revising the organizational model for RSSAC, and we're building a two-layer model where we have an executive committee, which is a rather small group which is focused on driving the process. And we are aiming at a larger group that we refer to as the "caucus," which will be a loosely-connected group with experts that we can reach out to to have help with actually producing documents and doing footwork.

So the executive committee consists of one voting representative from each of the root server operators, so that means 12 in total. And it also

consists of liaisons to various groups and other bodies of ICANN. The number 12 here stems from the fact that Verisign operates two letters. There are 13 root server identities but only 12 organizations that operate them.

We intend to create the initial process and procedures for how RSSAC is going to work in this new shape. And we also, in the executive committee, are going to select and keep track of the work items that we deal with. We're going to appoint work parties that deal with the various issues and draft documents by selecting subgroups from the caucus. And then the executive is responsible for publishing the results of that. I'll go a bit more into the details of the publishing process in a minute.

We also appoint outgoing liaisons. This is currently limited to the Board and the NomCom. And we also accept incoming liaisons from groups that want to have liaisons with us. And I have a list of those further down the line.

And also, of course, we elect two co-Chairs to run the operation. So the current two co-Chairs are Professor Jun Murai from the WIDE project in Japan, representing M-root, and myself, from Stockholm, Sweden, representing I-root.

And the current liaisons we have are from outgoing to the ICANN Board and NomCom, as I mentioned, and incoming from the IANA, from the NTIA, and from the root zone maintainer, which is Verisign. So we have liaisons from all the parties that participate in the root zone administration. And we also have incoming liaisons from the IAB (the Internet Architecture Board) for protocol-related issues, from the SSAC

for security and stability related issues, and also from the Government Advisory Committee.

The caucus, then: it's currently not formed yet, but the intent is to have it made up from people with various expertise, from DNS protocol experts, DNS operations people who operate DNS servers – actually on both sides, I hope, both on the authoritative side where we come into the play, where the TLD registries come in, but also on the resolver side where the Internet service providers and the end-users sit – because they probably have expectations on how the root service is about to perform, and hitherto, they've been a bit underrepresented.

The caucus will be called upon to form work parties, to subgroups, to work with the various issues and produce draft documents. But the caucus as a whole is also expected to review all documents, meaning documents from the other work parties that you may not participate in yourself. And the intent is to have this consensus-driven so that we can find things that we have consensus around.

The current status on this new group is that we are currently developing operational procedures for RSSAC regarding elections, liaisons, work party formation, publication of documents, and the entire thing. We are going to have an editing session on that tomorrow, where we'll sit down and try to bring that work a bit forward.

And we're also going to have a working meeting tomorrow with the RSSAC executive, where we hope to nail down a work plan with some time limits in it, which we will make public when it's decided upon.

The publication process: the current plan in the draft procedures document that we're working with, it looks like this: the executive committee is the contact point, the focus for work items, so either if someone wants us to look at something or if someone within RSSAC realizes that there is work that needs to be done, the executive committee will keep track of those issues and prioritize and so on.

Once a work item is identified, a subgroup of the caucus will be established and will make a timeline and try to specify what work needs to be done and so on. The work party will then produce the draft document, which is then circulated within the entire caucus to try to find consensus. Once consensus is reached, the executive committee will publish the document in the RSSAC series of documents, which is hitherto rather thin but we hope to improve on that in the future.

The next steps we're looking at is to finish the initial procedures document. We're in a bit of a chicken-and-egg problem. We need to define procedures for designing this caucus. And when we've done that, we have the tools – the caucus being part of the tools – to devise documents. So we really need the caucus to finalize the document, but we can't elect a caucus unless we have the procedures for it.

So my current take is that we will need to publish an initial version of this procedures document and then we will appoint the caucus, and then we will ask it to refine that document and endorse it or endorse a new version of it – so appointing the caucus, of course, a very important next step.

We have inherited two documents from the old RSSAC, which are now sitting in process-limbo. They're pretty much ready-to-go, but we don't

have the process for publishing documents right now. And we'll talk a little bit about one of them and we will talk about the other one.

We also realize that we need to establish better relations with the other bodies within ICANN – not only the Board and the NomCom, but the other advisory committees and also the supporting organizations. So one of the reasons for being at these conferences is that we want reach out to you and we want to be approachable.

And of course, this is to pick up issues that relate to the root servers that are handling other parts of ICANN. Where sometimes – not only within ICANN but in many circles – people don't always realize that there may be secondary or tertiary effects on the root service, but we who actually operate them and deal with them on a daily basis have the experience to see that that relationship maybe exists.

And by participating more and being more involved, we hope to be able to advise on root service-related stuff early on in the development processes in the other bodies so that we don't run into the head-on clash at the end when everything seems to be a done deal.

I'll continue with RSSAC 001, which is not yet published but it's sitting there on the shelf waiting to be published. There was an old document published by the Internet Engineering Task Force (the IETF) called RFC 2870 – and it actually had predecessor as well, called 2010 – which specify a number of requirements of root servers and root servers operation. That document has been outdated. It's fairly old. It's probably ten years old by now, and it's outdated and needs to be updated.

When we started to look at that, we realized that that document actually contains two components. It has a protocol side to it, where you have expectations of how the root server is going to behave, protocol-wise. And it has an operational side to it, where it talks about capacity and operational-type things. And we realized that the IETF is not the correct body to specify – not the operational side of it, at least. And the protocol side, it's questionable whether the IETF as a protocol-developing body should be the ones who specify the protocol expectations of an operational service.

So we have ended up with splitting up this document into two, where one document specifies the protocol side of it and is not published by the RSSAC and that the second document purely focuses on the operational side and is published by RSSAC, because RSSAC is the only focal point for the operations of the root servers, where people get together and can form an opinion on advice on this.

So the RSSAC document specifies expectations on a number of points, including infrastructure, service accuracy, service availability, capacity, operational security, diversity of implementation, monitoring, and communication. And all of those are dealt with in this document.

And it goes hand-in-hand with a document that is published by the Internet Architecture Board – or about to be, because we need to publish these in-tandem, connected to each other, because they reference each other. You should be able to follow the references when you read the documents, so they need to be published at the same time. So that document is also waiting for publication.

That document contains the protocol side: what should you expect of how the server behaves when you ask it a certain query? And that's a very short document, actually, which says more or less it should adhere to DNS standards as specified in RFCs and a few more things regarding IPv4 and IPv6 and so on. I think that's it.

The next document that we also inherited is the RSSAC 002, but I would like to ask Bill to speak about that a bit.

BILL MANNING: Okay, so I will talk to that corner of the room because that corner of the room is empty. So the origin of RSSAC 002 was—predates back to March 2011. In 2011, there was a meeting of some folks at the San Francisco ICANN meeting, where we agreed to try and document what needed to be measured, because if we're going to have a system globally that is responsive, there needs to be some sort of empirical or third-party ability to measure.

And so we agreed on these particular characteristics. The latency of the distribution system: from the time an authoritative copy of the root zone is created until the time that that data is available to be queried by the Internet community at large. How long does it take to get a consistent zone file to all of the root servers?

The size of the overall zone: there were some fears that going from 300 to 3,000 or 300,000 or 3 million entries in the root zone might affect the propagation and the ability of the root servers to actually pick the information up. So, "How much data is in the root?" was a question.

The number of queries received: how many questions are asked of the root server system? What is the trend on the number of queries? Is it going up; is it going down per server? All those things.

The query and response-size distribution: so when you ask a question, the root servers will give you an answer. Historically, the answers were relatively small. With the addition of DNSSEC and IPv6, those answers can be quite large and can be used as a denial of service attack vector. So we want to be able to measure that to see if we are, in fact, being used inappropriately.

And then an RCODE distribution, which is DNS-speak for, what kinds of questions are being asked?

And then the number of sources seen, which is really, who is asking the questions?

If you are uncomfortable with this list, please let us know and we will think about changing, augmenting, or deleting items on the lists to get the kinds of questions and answers that you want. But at least initially, this is the type of stuff that we're looking for.

Things that we are explicitly not measuring or going to be publishing – some of the operators will, in fact, collect this but it won't be published as an RSSAC activity – which are queries that are inappropriately or malformed or bad. Those types of things are outside the scope of RSSAC 002 measurement.

There are a set of concerns that have been expressed about this work. From an operational perspective, if we're going to collect all this data, we're going to have to backhaul the data from root servers around the

globe to be able to analyze it, collate it, and present a coherent picture. It's possible that there will be an effect on available bandwidth. So if we're answering queries, we may not be able to move the data back in a reasonable period of time.

And collecting that measurement data, some root server instances are not exceptionally overpowered. And so if we add additional load, we may stress them. And that would require upgrading equipment in the field or bandwidth. And it's not clear where the money for all that's going to come from.

The current DNS software logging limits that inhibit perfect collection and resolution of the latency due to the lack of serial numbers in the AXFR/IXFR logging statements – these are technical things which basically say in the DNS protocol itself that we use for doing the distribution, we don't have the tools to actually timestamp things.

Some additional concerns: the latency and distribution system could be more granular and also affect the time it takes for root name server instances to commence serving from the zone upon receiving it. In practical terms, the reporting feature is not currently available in the DNS software. So we are trying to do things that currently are difficult, if not impossible, to do in the current deployment strategies.

We are seeing that with the larger responses, that UDP is no longer always sufficient and we're seeing some move into the TCP protocol. TCP has its own sets of issues in the deployed global Internet. The fragments are a non-trivial exercise to capture and provide meaningful statistics. In some cases, we may think of them as bad or malformed

and not collect them at all. However, if we are seeing a transition to more TCP, we really do need to measure that stuff.

And that's the last one. So I guess the question at this point is: do you have questions that you would like to come up to the mic and ask? And I see we have a willing victim – two willing victims.

SUZANNE WOOLF: We should let our Chair manage our queue.

LARS-JOHAN LIMAN: Yeah.

SUZANNE WOOLF: It goes with the job, man.

LARS-JOHAN LIMAN: Roy, please.

ROY ADAMS: Hi. I'm Roy Adams. I work for Nominet. I noticed on, I think one of the very first slides, about service expectations. It basically specified a set of service expectations that root server operators must specify. Do you expect all root server operators to specify that sort of expectations? That's question number one. Question number two is: after the root servers have specified those expectations, do you expect the root servers to be held accountable for these expectations?

LARS-JOHAN LIMAN: To you first question, I had a bit of problem to interpret it. Did you say each root server should fulfill those expectations, right?

ROY ADAMS: No, but you said "must specify."

LARS-JOHAN LIMAN: But you say "specify," but how do you mean to "specify"? I see this document as the specification. Or do you mean to put more-detailed numbers on them? Or what do you mean by "specify"?

ROY ADAMS: Well, that's a good question. I read specifically from the slide: "A set of service expectations that root server operators must specify."

LARS-JOHAN LIMAN: Ah, sorry. That's probably a…

SUZANNE WOOLF: It's "satisfy."

LARS-JOHAN LIMAN: Yes, "satisfy" would be the term.

ROY ADAMS: Oh? It should be "satisfy"?

LARS-JOHAN LIMAN: Yes.

ROY ADAMS: Okay. I thought it said "specify." Good, then my question is almost the same. Do we expect all root server operators – all 12 – to satisfy those expectations and to comply with it, and are they measurable?

LARS-JOHAN LIMAN: I would actually, yes; maybe not from day one, but over time, yes. And different operators have different environments that they operate in, so there may be small differences. But overall, I would expect them to do that, yes. Bill?

BILL MANNING: As one of the root server operators, we have had some question about what these requests actually mean in the context of personally-identifiable information in some parts of the world. If it is properly anonymized and no PII is released, then there's no question that this data will be collected and made available.

What's unclear is whether there is PII here that we have to be concerned about. Even if we're serving the root from one jurisdiction, if it shows up in another where that is a problem, we're still liable and so there's a concern there.

ROY ADAMS: I understand. I think you're answering a question that I did not ask.

BILL MANNING: The answer is if there's a legal problem with collecting and publishing the data, we won't do it.

ROY ADAMS: I was referring specifically to the service expectations. Okay, so yes, I'm talking about this slide, specifically. I mean, RFC 2870 is going to be replaced, obviously, and hopefully that's released through the IAB or through the IETF. The new RFC replaces the old one, etc.

BILL MANNING: Ah. Okay, so this particular question, the answer is that 2870 and its predecessor, 2010, talked about not the root server system but root servers themselves. And to the extent that RSSAC 001 talks about the hardware platforms and what the expectations are about the hardware platform, those are less-material than the overall performance of the system as a whole. And so asking about a specific operator or asking about a specific platform becomes moot if the system is well-behaved.

ROY ADAMS: I disagree. I completely disagree.

SUZANNE WOOLF: Okay, let me dive in. I love this question. Roy, you've made me very happy. And none of us is speaking for the committee overall, but I think only because we don't have a formal consensus mechanism at this point. I think any of us is comfortable speaking of the state of play as root server operators and as members of the committee.

Two things I would say to you that I do think are responsive to your question: first of all, part of what has taken time and effort over this document is working toward something that does have consensus support from the root server operators in the sense that they will commit to observing the service levels and commit to publishing that data that will allow people to see that we're meeting those service levels.

And these other questions about exactly what's published and exactly what the statistics mean and exactly what the PII involvement might be and so on, those are details. And I think the implementation question is is there consensus behind this document to meet these standards and to publish data that will allow people to observe that we are? And I think the answer to you on that is yes.

The other piece of my answer to you is that all of these things we're saying about engaging more broadly with the community and having mechanisms for input from other portions of the community are exactly so that we can evolve these expectations in collaboration with the community and [send text].

LARS-JOHAN LIMAN: [And Roy], thank you very much.

ROY ADAMS: Thank you.

DANNY MCPHERSON: Hi. You had some slides up there: What We Do. Can you go back to that?

SUZANNE WOOLF: It will take a second.

DANNY MCPHERSON: No problem. There's two slides back-to-back, yeah. The What We Do and What We—that one first.

So we're talking about RSSAC here, right. So you don't actually do the first thing there, right? Okay, so if you go on to the next slide, it says, "What we don't do is we don't tell root server operators what to do." Who does?

LARS-JOHAN LIMAN: I would say that the root server operators, they listen to community through various channels. RSSAC is one of them, but there is no formal process today for telling a root server operator what to do. We're trying to put in place various checks and balances to remedy that situation. These RSSAC documents is part of that. But currently, there is no formal process for telling a root server operator what to do.

DANNY MCPHERSON: We have an RSSAC document that is going to make some recommendations, presumably to the root operators, that the root operators can completely ignore. Is that what I heard?

LARS-JOHAN LIMAN: They can completely ignore, probably yes. But I would say that it's very unlikely if they are well-founded recommendations that have a meaningful interpretation for the root server operators. I would expect them to listen carefully to that.

BILL MANNING: So backing up just a little bit, about four years ago, the question came up as to what will the roots do? Can the roots arbitrarily manipulate the contents of the root zone and get away with it? And the answer turned out to be, yeah, we probably could.

But each of the operators was asked to consider making a statement or having either a unilateral or bilateral agreement with the IANA functions operator that we would only publish the data that they provide unaltered in a timely manner. And most of us, as operators, actually made such a statement.

So our good word that we're not going to screw up is basically our good word. And the attempt at this point is to actually find better ways to become more accountable to the consumers of the data that we publish. And I don't quite know how to do that. So if you can help us – and you too – figure out how do we become accountable to the people that use the data we publish, that would be great.

DANNY MCPHERSON: So you want to clarify something, and then I'll go on?

PATRIK FÄLSTRÖM: Yeah. I would like to add another clarification, which made things a little bit complicated.

LARS-JOHAN LIMAN: And who are you?

SUZANNE WOOLF: Who are you?

PATRIK FÄLSTRÖM: I'm coming to that. Patrik Fälström. I'm working for Netnod, also I-root. I also happen to be chair of SSAC, but that doesn't matter. I am currently a root server operator.

Netnod is based in Sweden and Europe, and in Europe we have a directive for telecommunication that very clearly say that for providers of electronic communication services, if it is the case that there are disruptions, that should be reported to the responsible regulator.

Now, the question is then: is the root server service that Netnod provides, does that fall under that telecommunication regulation? And their answer is maybe, and this is something that is currently under discussion. So it might actually be the case that for the actual service that is provided, that it falls under whatever telecommunication regulation that exists in the various countries.

Now, that is something that, of course, only says if the service fails. The question is then: what does that mean? So we also have to be careful

when we're asking the question, what part of the question we are talking about.

BILL MANNING: That's why we need RSSAC 002, so that we actually have some measurements so we can come back and say, "The system is healthy," or, "The system is failing." We need the measurements.

LARS-JOHAN LIMAN: Suzanne?

SUZANNE WOOLF: Yeah. And in fact, two things about that: one is that people do publish a variety of measurements, and the document under discussion is about publishing a mutually consistent set so that the results can be examined across the system or in any individual operation.

I think the other thing that's germane to your question, though, on the specific issue of being able to tell if the data that you got out of a root server is what the root zone maintainer put in, DNSSEC has given us the ability to authenticate that data. And that's an important piece of the puzzle; you're no longer just trusting whoever is answering your root name servers, your root queries. You can validate.

And most in practice, we know that DNSSEC validation is not yet widespread, but the capability is there and anyone who wants to can look and validate the integrity of any answer they get from any root server any time.

DANNY MCPHERSON: Yeah. No, I certainly agree with that. That wasn't my question at all. I think that given that we have DNSSEC, anyone could serve the root zone any manner they wanted and there's not really much other folks could do about that. I was actually talking about the root server operators that may or may not have some relationship with ICANN or the IANA function.

SUZANNE WOOLF: Okay. See, I thought it was germane to the wider question, but thank you.

DANNY MCPHERSON: Yeah, it's important actually because now it doesn't matter about the issue of a potential malicious root operator because relying parties can validate the contents and object level integrity of zone and all that. Yeah, I got it.

I was more concerned with things like SSAC recommendations that were made and things like name collisions where we could measure the impact of things like community outreach. And we could say, "Hey, here's the 1,400 applied-for strings, and here's the measurement apparatus of the root system, and here's the volume and type of queries and the distribution."

And the current trajectory I see, that's never going to happen. And I think that's really unfortunate. I think it's actually gloriously ironic that everyone seems to be concerned with an NTIA/IANA transition, and yet

this room is mostly root operators. And yet this is the operational system that we're talking about today.

And so I think that, yeah, there's certainly room for change in this. I think that I certainly would like to see the publication of these and in the documents and the procedures that RSSAC publishes beyond their engagement with the Board or with other aspects of the community, the accountability models they hold one another to.

Because I think there's a glorious lack of accountability across the root operators, and I think that if you're going to be a root operator, that you should agree to a certain minimum set of standards that the root operators and RSSAC and the community think are sufficient and then you move forward. Otherwise, as you stated with DNSSEC and object-level integrity, anybody can distribute a sign zone and relying parties can choose to validate it or not, and they can just ignore the global root. And I think that would be really unfortunate.

LARS-JOHAN LIMAN: Thank you. Next and last, please, because we are running out of time here.

[UNIDENTIFIED MALE]: Okay, it's [inaudible] from the [inaudible]. I'm not a root server operator. [inaudible] thank you [inaudible] root server operators for their contribution in the past many years. I'm very happy to hear that RSSAC has planned to be very organized and to be much more open.

Regarding to the RSSAC, there one [inaudible] slide, I think now the [inaudible] service expectation. I'm not sure if that is [inaudible] or the similar as the service level agreement.

LARS-JOHAN LIMAN: Pardon.

UNIDENTIFIED MALE: I think from this recommendation we use to [inaudible] for service expectations. So I just want to confirm if that is similar with the service-level agreement. So does that mean that the root server operators will promise to have some kind of service level for root server operating?

BILL MANNING: Service level agreement with who?

UNIDENTIFIED MALE: With the community.

BILL MANNING: Great. Let's talk after this meeting.

UNIDENTIFIED MALE: Okay. So if I read the RFC 2870, [inaudible] and the requirement is to [learn]. And for the RSSAC, there are two [inaudible] for the moment. I think maybe you can have some kind of [measurement] for the root server [inaudible] distribution. So the idea now there is over maybe

about less than 400 [inaudible] around the world. But I don't think the distribution is perfect, so if you need help, I can give some support.

BILL MANNING: I think that the first attribute in RSSAC 002 was measuring the latency of the distribution system, which might talk to your concern.

UNIDENTIFIED MALE: Okay. And also second suggestion is for the root servers Internet website is too slow to access and [inaudible], I suggesting have a mirror site.

LARS-JOHAN LIMAN: That's interesting. I would like to know more about that.

Thank you. I'm sorry. I did close the queue after you because we're running out of time. We only had until a quarter past, and we are already running out of that. I'm happy to talk to both Danny and [Hans Peter] after the meeting. Please come and see me, but we really need to close this meeting here and now.

So thank you all for attending the meeting, and please do come and talk to us. We're open. We want to talk to you.

SUZANNE WOOLF: Yes. There's lots more dialogue still to have.

BILL MANNING: Yep.

LARS-JOHAN LIMAN: Thank you.

BILL MANNING: Thank you.

[END OF TRANSCRIPTION]