
MARRAKECH – RSSAC Public Session
Wednesday, March 09, 2016 – 14:00 to 15:30 WET
ICANN55 | Marrakech, Morocco

UNIDENTIFIED MALE: Good afternoon, everybody. We've got lots of room at the table if anybody wants to come up, join us. We won't put you on the spot. I'm Brad Verd. I am the Co-Chair of RSSAC. And my colleague, Tripti Sinha, was unable to make the meeting. She sends her apologies, and looks forward to seeing everybody at ICANN 56.

Let's jump into the slides. We'll try to get through this as quickly as we can. Here's kind of a quick overview of the agenda today. We're just going to run through some documents. Server expectations of the root, measurement stock, RSSAC 002. We'll cover the workshop report, do some quick updates as to what's going on right now, and end with kind of a Q&A, and membership run through.

As I said, Brad Verd, this is the reoccurring message that comes from RSSAC. Trying to educate everybody so that they can recite it for us. The role of RSSAC, the Root Server System Advisory Committee is very narrow in scope, but key to the DNS ecosystem. Really quick, we are here to advise the ICANN

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

community and board on matters relating to the operation, administration, security, and integrity of the root server system. Again, very narrow in scope.

This is RSSAC the organization. The formal committee, the formal RSSAC committee is made up of representatives from each of the 12 root server operators. Each of these have an alternate, and we have a number of liaisons from our root zone management partners, as well as from the ICANN community. More importantly, the caucus, which is the body that does all of the work, and provides the output of RSSAC, these are made up of volunteers, subject matter experts. Right now, we're somewhere beyond 70 members, and a large portion of those are not related or directly related to root server operations.

Each of the members, in our goal to be fully transparent, they provide a statement of interest that is available on the ICANN website. Again, the caucus gets all the credit for the output of our documents that are published. Purpose, pretty simple. Our goal is to produce documents that are based upon questions and concerns from the community and from the board. Again, our goal is all about transparency. We want to be as transparent as possible in everything we do. And we've defined our framework to get that done. There is our e-mail for our RSSAC membership, if you want to join. Please state your interest to that e-mail address, and we'll get back to you on that.

Publications since ICANN 54 in Dublin. RSSAC 001 has been completed and published. This document has been completed for a while. It's been sitting on the shelf. As we were trying to get the tandem document from the IAB done, which you can see is RFC 7720. These documents together move the legacy document, or the old document of 2870 BIS into a legacy state. The IAB document defines the protocols for the root server operators, and RSSAC 001 defines the service expectation of the root server operators and servicing those protocols.

002 version 2 has been published. 002 is the measurements document, where it defines a number of measurements and metrics that each of the root server operators are supposed to publish and make available to the community. This is the second version of that document. There were some minor updates that were made based upon real world experience and things that kinda needed to be tweaked a bit to make them work.

In addition, RSSAC had two statements that were published. One on the RSSAC workshop, so this was the first of its kind for RSSAC. We gathered the formal committee, locked them in a room for a number of days, and talked about root server topics. There's a report there that is on what came out of that, and I think we're going to cover that later today also in more detail. In addition, there was just a public statement on the CCWG accountability proposal that was done.

I'll go into a little more detail on 001, just to refresh everybody's memory. As I said, it's been done for a while, but it is now published and official. This document establishes a set of expectations for root server operators, how they're supposed to run their service. It defines 18 operational expectations, which I'll go into those in a bit here. As I stated earlier, together with RFC 7720, it makes RFC 2870 legacy.

These are the different areas that are defined in 001. Pretty much everything you can imagine around the operation of a root server operator. And most of these, each root server operator is supposed to publish something related to each of these topics, and make it available on the root servers webpage, www.root-servers.org which will come up here in just a bit.

Here are some of the things. Each root server operator is to publish the level of service that they offer, and that's just to kind of set expectations with the community. The topic here is where everything is deployed. Again, all of this is on www.root-servers.org website, and is update on a constant basis. It is current. I'm going to turn the mic over to [inaudible] from K-root to cover the RSSAC 002.

UNIDENTIFIED MALE: Good afternoon. My name is [inaudible] from Ripe NCC, and I'm in charge of K-root. Yes, RSSAC 002, basically this is almost all of

or many of the root operators already published statistics on what they do. For example, at K-root for years and years, we published the amount of bandwidth which were used by each [inaudible] and query loads, and things like that. It was never a consistent thing between all operators, and it was never the same format based on the same definition. So RSSAC 002 tries to address that. It basically identifies and [inaudible] set of measurements, to establish a base line and have a trend.

And having a trend is very important, especially nowadays with all the changes. And the size of root zone from the publishing side, or TTL changes, or to the consumer side where you have IPv6 deployment, DNSsec, queries over TCP, all of these variables. So having a baseline, and being able to look at the trends helps a lot. And it also helps to establish an [inaudible] running system. Because, based on this information, you can actually see if something is going wrong much more earlier than finding out after something has gone wrong. These were the two, basically, main requirements for having [inaudible].

The document, basically, defines six high-level metrics. They're listed [inaudible] in publishing [inaudible] data, size of overall root zone, number of queries, query and response-size distribution, R-code distribution, number of sources which are seen.

Also, the document [inaudible] have a few recommendations. For example, it suggests that each root operator implement the measurements in the advisory. Every single letter has acknowledged that, and they're approved to do so. It also states that RSSAC should monitor the progress of the implementation, which RSSAC is doing, and I will report back on that. It mentions that the document should be revised, a maximum in two years from the date of publication to adopt the new technologies, and changes in usages, and things like that.

This is the implementation status. As you can see, 9 out of 13 are already publishing. Out of four, you see the expected completion time, which is not that far in the future. So basically, well before the end of this year, we will have this report from all of the letters. How to get to the statistics? If you go to www.root-servers.org, and click on any of the nine letters that are already publish RSSAC statistics, you can basically see bottom, which is surprisingly called RSSAC. If you click there, you will see a set of five which are published daily. Per day, you have six files, basically. For each high-level metric you have one file per day.

For example, there's an example here. As you read, it says from A-root, [inaudible] it's for 1st of January, and this is the traffic volume metric. You can read that there are multiple metrics here presented. To see how they can be useful, for example, it says that there were 4 billion IPv4 UDPDNS queries 1st of January to A-

root where there were 200 million IPv6 queries. So it's about 5%. And you can even look at, depending on what you're interested in, you can look at the responses for example. So you had this many queries, this many responses. And it's immediately visible how useful this can be, especially if you establish a trend, because you can see the trend of queries, and how they change and all this.

The [inaudible] format, and the first version, after we started implementing some of the letters, we found out some parts are open to interpretation, some parts are not clear enough, and there are some [inaudible] syntax servers [inaudible] some indentations missing, or units that were not clear, or some of the metrics were not clear. That was with a very good work of [inaudible]. That was actually clarified. Everybody agreed on one set of standards and definitions, and that's as presented, that's now in place. The version three of the document is upcoming, as my colleague [inaudible] will present in the upcoming caucus work.

LARS-JOHAN LIMAN:

The next slide is me, but that's not what I'm going to present. My name's Lars Liman. I work for Netnod. We operate I-root. I'm going to tell you a bit about the workshop that we held last fall. As Brad mentioned, we decided to get together and actually

work in a concentrate way for two days straight. That was actually very useful, and we had lots of rewarding discussions. We had a couple of goals for this workshop. It was to discuss the evolution of the root server system. And we wanted to come to a common understanding of what the current platform actually is because there are lots of things that are hidden in folklore, and people's memory, and so on. And we needed to discuss, and come to a common understanding of what the current platform is. We made good progress on that.

We discussed evolution, accountability, and continuity of the root server system. We came to a consensus on a number of points, and those have been documented in the workshop report, which I believe has been published. So it's somewhere on the RSSAC pages. How you find the RSSAC pages is a different thing, not our fault, use Google.

If we look at the actual consensus points that we reached, we divided them into accountability, continuity, and evolution. Not necessarily in that order. On the accountability side, RSSAC as a whole acknowledges that root server operators have responsibility to the Internet for providing good overall service.

This may sound like a very obvious point, but it's anything that we actually want to [inaudible] in text somewhere. We also understand that in addition to the above, there is also various

forms of accountability to local communities. For instance, in our case, in Netnod case, we have a responsibility towards [inaudible] the academic network in the Nordic region because we have a memorandum of understanding with them. RIPE NCC may have relations to their membership. Various root server operators have local communities, or supporting organizations, or something that to which they have accountability.

We also identify that the root server operators actually have responsibilities towards each other. We actually have to cooperate. We have to operate a combined system that works as a whole. We cannot have one or two operators go out on their own and do their own thing. We have to operate as one common system, and that creates responsibilities between us.

We also wanted to put in text something that has been obvious for a long time, but we really wanted to have that written down, which is that the RSSAC is not aware of an existing procedure to designate a root server operator. There were procedures that were extremely informal. They vanished in 1998. Since then, we are not aware of any procedures to designate a root server operator. We also suggest that the process to do that be defined because it may be needed in the future. It would be good to have one, and RSSAC would expect to participate in the development of such a process. That's not codified here, but it's obvious that it's a community effort that needs wide participation.

On the continuity side, we said that RSSAC should document the position, including the underlying reasons for having that position, which is that the demise of what a single root server operator does not pose an immediate problem for the root server system as a whole. Because the system is so over provisioned in all possible ways, both when it comes to technology, and organizations, and what have you. If one of the root server operators suddenly stopped operating their letter, that wouldn't have any noticeable impact on the system.

We also should document that the various aspects of diversity, which are many, of the root server system and operators, these are strengths of the overall system. It's actually beneficial for the entire system to have a lot of differences in many places because that creates a web that is very durable. The root server operators should also document and promote the long track record of uninterrupted service operation, and the operation of the root server system in order to just highlight that this is actually a system that has been working for a very long time.

If we go into evolution, we see it as desirable to define the key technical elements of potential root server operators – the technical elements that would be critical to define a root server operator – sorry, I have to get this right. Yes, what we're trying to say here, is that we need to define critical elements that can form the basis for a process for designating a root server

operator. Because where we work and the important things to us are the technical components, and how that fits into the entire picture. This is a rather big task, and it's something that will require work on a lot of planes to define this. But we see that this is a very important starting point for creating the process of designating a new root server operator.

We also recognize that the root server operators need to stay ahead of the demands on capacity and performance. We cannot just sit back and say, "Oops it broke. It's time to fix it." We must be ahead because we cannot afford that this component of the Internet as a whole actually stops working.

We've also realized that RSSAC is not a poster child for transparency and communication with other groups. So we need to create a strategy for better transparency and better communication with other parts of ICANN, with the Internet community overall, and so on. That work is already undergoing.

This last bullet is actually a very important point. That emerging technologies affecting the root server system should be embraced, as long as – and this is the very important part – as long as the Internet globally unique public name space is preserved. We maintain an openness towards new technologies that can involve, affect, replace, whatever, the root server system as a whole. We cannot really sit back and say, "No, we

must not ever change.” Because change always happens. We must be open to that. The key important thing here is that we maintain a unique name space, because if we don’t have that, we no longer have a functioning Internet, in my personal view.

Handing it over to Susanne Brown. I don’t know who is going to do this.

UNIDENTIFIED MALE:

Susanne couldn’t make it here today. Susanne Woolf is our RSSAC liaison to the board, and there is a board meeting going on that we all agreed she should be at. You guys get to listen to me for a bit more.

The current work, we have a number of things that are happening that we expect to get done before ICANN 56. The first one here is the history of the root server system. Going through our different documents, and our workshop, and different efforts, it became abundantly clear that there was no document that could [inaudible] or memorialized the history of the root server system. So we thought that this was an important fact that needed to be put together, and lots of people worked on this.

We worked in collaboration with root server operators both past and present, as well as a number of different technologists who

were involved, dating back to the beginning of the root, and put together this document. That, and now with the help of ICANN staff, thank you very much, this has come together very, very nicely. It contains a chronological history going back to the beginning as well as all the current operators have added information about their organizations, their roots, and major milestones that they may have done independently.

Before I go on there, that's a great document. It's a great read. It's currently being sent around with the caucus – or no, actually, I'm sorry, it's out of the caucus, and now we have some final edits going on, so hopefully that'll be out soon for the consumption of the community. RSSEC 002 version 3, so you heard [inaudible] earlier describe the measurements document, and what was going on there. This is version 3 of it that will have some significant changes or additions, I should say, on the stuff that we currently measure, and adding more.

Lastly, we have a work party that's going on regarding the root server naming system scheme. This is as the Internet have evolved, or the root zone has evolved, different changes have taken place to accommodate the evolution or the ever changing root server system. We thought that it was time to look at this again. Have we named the root servers correctly to get the most out of the environment we are currently in? Some of this, as I said earlier, started the discussion around the history document,

and that was part of these discussions. The work party is supposed to look at new naming schemes, do a risk and impact assessment, and then come back and give recommendations, if there are to be made any changes, what are those, and what is the scope of those, and everything related to those potential changes.

That covers work that has been done since last ICANN, work that is currently ongoing that we hope and expect to be done before ICANN 56. And we are moving into our Q&A portion now. If you go to the next slides. This is community interaction. This is where all of our resources are. Please feel free to go there. I know [inaudible] mentioned Google, but right here are the links. Again, there is an e-mail to our membership committee, if you are interested in joining RSSAC and contributing. I'll open the floor up to questions.

All right, I think with no questions, I think we'll give you guys some time back to go enjoy the nice weather. Thank you.

[END OF TRANSCRIPTION]