
PANAMA – RSSAC Information Session
Thursday, June 28, 2018 – 09:00 to 10:15 EST
ICANN62 | Panama City, Panama

TRIPTI SINHA: Alright, let’s get started. Welcome, everyone. This is the RSSAC open session information session and we’ll give you a quick one-hour synopsis on what RSSAC is all about and what our activities are. We’d like to start with introductions. Let’s start with Kaveh. Are you the last one at the table? If you could introduce yourself, please.

KAVEH RANJBAR: Kaveh Ranjbar, RIPE NCC.

JEFF OSBORN: Jeff Osborn, ISC.

FRED BAKER: Fred Baker, also ISC.

TERRY MANDERSON: Terry Manderson, ICANN.

TRIPTI SINHA: Tripti Sinha, University of Maryland and co-chair of RSSAC.

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.

DUANE WESSELS: Duane Wessels from Verisign representing the root zone maintainer.

BRAD VERD: Brad Verd, Verisign, co-chair.

LARS-JOHAN LIMAN: Lars-Johan Liman, Netnod.

RYAN STEPHENSON: Ryan Stephenson, Department of Defense.

RUSS MUNDY: Russ Mundy, SSAC liaison to the RSSAC.

WES HARDAKER: Wes Hardaker, USC ISI.

KEN RENARD: Ken Renard, representative Omni Research Lab.

DANIEL MIGAULT: Daniel Migault, IAB, SSAC liaison.

TRIPTI SINHA:

Alright, thank you. Next slide, please. The way we've split the agenda today is we'll give you a brief overview of RSSAC and what it's all about, talk about our publications since ICANN 61 and then an update on current RSSAC and caucus work and then interaction with those of you in the audience. Next slide, please.

Essentially, the Root Server System Advisory Committee is here to advise the ICANN board on matters relating to the administration, the operation, the security, and integrity of the DNS root server system. It's a very well-defined scope. Next slide, please.

We are all, each and every one of us, appointed representatives from the 12 root server operators. We have primaries. We're represented by primaries and alternates. We also have liaisons, outgoing and incoming, depending on what is required to help assist in our work.

We have an extension to ourselves called the RSSAC caucus. Essentially, they are a body of subject matter experts and they are appointed by RSSAC. Next slide, please.

To date, we have 92 experts, DNS and root server system experts, in the caucus and each of them has a public statement of interest online and they do receive credit for individual work.

The purpose of the caucus, as I said earlier, is it's a pool of experts with a certain amount of expertise and it's quite a broad set of skill sets. There exists transparency of who does what work and what expertise they bring to the work. It's a framework for getting work done.

In other words, we in essence, define potentially a certain result that we'd like to see with a problem statement. We define leaders and also have deadlines to the work.

If you're interested in joining the RSSAC, the caucus, please do send an e-mail to rssac-membership@icann.org. We are meeting next at IETF 102 in Montreal, Canada. Subsequent to that, at ICANN 62 in Barcelona, Spain. Next slide.

This is a very quick update on our publications since ICANN 61. RSSAC 38 is an advisory that we presented to the ICANN board on potential governance model for the root DNS system, and 37 is the actual meat of that proposal and I encourage all of you to look at that piece of work.

RSSAC 36 is a statement on the draft final report of the organizational review that RSSAC has recently undergone and it was our response.

35 is a final draft report of the second organizational review of the Nominating Committee, so our comment on that Next slide, please.

34 is essentially a report we do. We conduct workshops and that's a synopsis of what we did in our workshop in May 2018.

RSSAC 33 essentially gives you a definition of what RSSAC and Root Ops happen to be. RSSAC is essentially an advisory committee. It's an artifact of ICANN, the community, the organization. Root Ops live outside of this construct and is a loosely coupled organization where the operators meet. I encourage you to read RSSAC 33 if you'd like to know the difference between the two.

32, once again, ties right back to the review that we recently underwent. The way the reviews are split up, they do an assessment and then recommendations and that was our feedback on the assessment. Next slide, please.

I thought we'd very briefly give you an update on what 37 is about. That's the big piece of work that we just ratified. It's a governance model for the DNS root server system. What you see, the diagram that you see, is a very high-level abstraction of the proposed model. It instills 11 principles that we believe should be maintained in the model going forward.

It proposes ... It is an initial governance model. It's a starting point. We demonstrate how the model works through a set of scenarios on how these different situations could occur, how this model would operate if you were to designate and remove operators, and that's contained in the 50-page document. I highly encourage you to read it. It also talks to how this would be governed in the future. Next slide, please.

RSSAC 38 is the actual advisory that tells the board, "This is our advice to you." Initially, we say we have produced a model called RSSAC 37. It is simply a starting point and we encourage the board to put some structure around it, to implement it. I open it up to the community to get more feedback. Also, step two, we tell the board to initiate the estimation of the cost of deploying this model and implementing it after the community has arrived at a final version. Next slide, please.

33 is essentially a statement, as I said earlier, on the distinction between the RSSAC and Root Ops. There's been some confusion for many years on the two different concepts and bodies. Let me go down through column one, which tells you what RSSAC and the caucus happen to be and what Root Ops is, in the right-hand column.

RSSAC and the caucus. RSSAC is essentially an appointed body created by the ICANN board and it's here to advise the ICANN

board and the communities on matters related to the operation, the administration, the security and integrity of the RSS. We generally devise [inaudible] comment statements, procedures, and reports. We also collaborate with the RSSAC caucus and we come up with statements of work and organize and shepherd the work that is assigned to the caucus. And we engage with the ICANN community regarding issues pertaining to the root server operators.

Some RSO staff members participate in the RSSAC and the RSSAC caucus but it may or may not be staff members that have any operational responsibilities. So, that, essentially, is the RSSAC and the RSSAC caucus.

On the right-hand column, we describe what Root Ops is and [inaudible] actual technical operation of the root server operators. Essentially, we implement technical advice from RSSAC and IETF in other forums, and it's the day-to-day management and operations of the infrastructure that makes the DNS root zone available to the client ecosystem of the Internet and based of course on RFCs.

It's a forum where the RSOs collaborate with each other to ensure that the integrity and reliability of the RSS is maintained and discuss some very critical issues regarding operations and the security of the system.

It also shares any kind of DNS statistical data with the Internet community. Some of the staff members participate both in Root Ops as well as RSSAC and the caucus in different roles. So, some of us do wear multiple hats. Next slide, please.

The other publication here is RSSAC 36 and this is our statement on the final report of the second organizational review. In this particular report, we do express some very serious concerns about the scope of the review as we do believe that there was a tremendous scope creep in this particular effort, and we believe they missed the mark by focusing more on community opinions and not producing any kind of evidence or attribution in where the data was found. It was a lot of community conversations, very hard to ... With a lack of tangible data that they could point to. Also, it was not an organizational review. Typically, an organizational review is you do an assessment of procedures, documents, and minutes and structure, form and structure, of the organization and none of that was done.

RSSAC 34 is a report from the May workshop. As I said, we conduct a lot of our internal work through workshops. It's a very high-level summary of what we did in May. Next slide, please.

Alright, I'm going to turn it over to my colleagues now. Liman, is that you?

LARS LIMAN:

Yes. Lars Liman here. One of the current work items that we're working on through the caucus – so, this is work that's mainly going on in the caucus – is a work party to investigate the harmonization of anonymization procedures for data collecting.

When DNS data is collected, you typically collect the DNS queries coming in on the Internet connection. They contain by necessity the source IP address of the entity sending the query to the server. And if you want to collect this information and store it and process it for technical analysis or statistical purposes, the data, the IP address, is actually part of the data that you process.

This is seen as somewhat [inaudible] an integrity, personal integrity perspective, especially in the new GDPR context. So, already today, some of the operators of large DNS servers – this primarily focuses on root servers, but it actually pertains to all types of data collection.

Already now some of the providers of data anonymize the IP address, so they can update them. They modify it, so it doesn't show the exact [inaudible] content. The DNS query remains unmodified, so you can actually follow what's being asked for, but not by whom.

This anonymization turns into a problem for the researchers because they cannot compare data from different operators. So,

this was a question to the caucus, a plea for help, to find out whether harmonization and anonymization procedures is something to recommend to the entire root server operator community, so that we all do it or where it's up to each operator to decide whether to do it or not.

If that question was answered by yes, to recommend a preferred way to anonymize the data so that we all use the same method, and then to consider whether to recommend that anonymization be undertaken by all who share data.

This work is very near completion. There is a draft report that has been handed over to RSSAC and we are in a dialogue with the work party in the caucus about adding a bit of motivation. It turns out, in their work with this document and their investigations, they've reached the conclusion that these questions weren't really the perfect questions to ask. So, in their work, they modified the work that we're doing. So, they provided very valuable input, but they didn't quite actually respond to the queries that we asked for and we would like them just to motivate that in the document, so that it's obvious why this work kind of shifted from the original instructions, which it was a good idea. It makes perfect sense in after thought, but it should be stated in the document, so that it matches the original requests. That's where we are at right now. I guess at the next ICANN meeting, this will be a published document. Thanks.

DUANE WESSELS:

Alright, thanks. My name is Duane Wessels. I'm the shepherd of the packet-size work party, also work currently going on in the caucus. The scope of this work party is to better understand all the intricacies of ways that the packets can be split up or the client is going to have trouble communicating with servers. There's a lot of complexity with large packets relating to MTU sizes, MSS sizes, ICMP, indicators of packets that need to be fragmented and so on.

Currently, there's a lot of diversity in the way the root servers handle this, which is generally taken to be a good thing. One of the goals of this work party is to understand if that diversity is useful in this particular case or if it might be better that there was more consistent handling of things like MTUs and MSS sizes.

Of course, further to complicate all this, things are slightly different in IPv4 and v6. We know that there are some really significant problems in the v6 space, so we look forward to this work party making some recommendations along those lines.

Current state is that it's sort of just gotten started writing its output documents. There is some software that has been shared and published that can be used to measure some of these things. Not all of these parameters, but some of them, can be measured with software. I expect that this work party will

continue for the next few months and have something, hopefully, by the end of the year.

WES HARDAKER:

[inaudible] about that. The RSSAC caucus has a GitHub repository. The purpose of it is to collect tools relating to the root server system, and additionally, the output of any tools that might be developed during an RSSAC caucus work party. That website is shown on the screen above which is github.com/rssac-caucus. Right now, there are five repositories available. One that was actually just added this week. So, there's an [inaudible] for accessing the RSSAC 002 data. The RSSAC 002 data collected from all of the roots is a repository in itself, so you don't have to go collect all the information yourself. You can just update it from GitHub in a mass collection, which is highly useful.

There is an anonymization harmonization repository that is the result of the work from the anonymization caucus working party, and there's root server naming tests, which is the work that derived from the rootservers.net naming and whether it should be signed and whether the naming system should be changed. That was a work party from a couple of years ago.

Then, finally, the most recent one added was the packet sizes working party which is just wrapping up sometime soon. The

results of that working party are also available via GitHub. Thank you.

BRAD VERD:

We in the caucus, we have collected all of our work items and we, unfortunately, have so many that we've had to kind of put them into a queue, and as work parties and resources become available, we will spin up new ones. But, this is a list of the work that is in our queue. We did a recent poll of all the caucus members to find out the high-interest topics and we are spinning up two new work parties on those, which are here. Here's the results of our poll and all the topics below. You can see that there were similar topics. We ended up combining them and we're working on two statement of works that will be going out to the caucus in the very near future. Can you go back to the original list? I'm sorry to jump around on you.

You can see that this was the full list that was brought up by caucus members, things that we should consider and maybe look at. Everything from resolvers to rogue root servers and what actual damage could they actually do. Geographic diversity, best practices for hosting. It was anything and everything that had been talked about in the caucus. Can we jump forward?

Obviously, we did a poll and came up with two items. The first one is we're researching the geographic diversity of the roots and going to study the behavior of how resolvers choose which root server to talk to. That statement of work should be going out very shortly to the caucus. Next slide. Wes?

WES HARDAKER:

Next we will talk about community interaction. RSSAC has, especially over the last few years, really tried to increase our level of transparency since the reboot a couple of years ago, actually. We're continuing that process and have made even more recent strides that we'll talk about in a minute.

All of our meeting minutes are published openly on the website available. If you search for RSSAC meeting minutes, you'll find the results of all of our meetings. All of our documents are published on our list of documents website which is searchable both by number and by date, for example.

It's worth noting that as we update documents, we let that document have the same number and attach a new version number to it, so sometimes if you search by date, you'll find that some of our older documents have been updated.

We routinely participate in the root server tutorial, so every ICANN – not the policy one like we're at now, but the longer two

meetings we hold two root server tutorials on how the root server works. If you have not been to one of those, you're encouraged to do so. There's always an opportunity to interact with a Q&A kind of session to learn more about the root server system.

Our operational procedures are published in RSSAC 000, currently version three. We're holding a conversation later today I think about what we might want to change in that document in the future, so it's possible version four might come out in the next couple of years.

Then, RSSAC meetings most importantly are now planned to be open by default. This is actually our most recent change. RSSAC meetings in the past have occasionally been closed due to conversations that we need to have in a private manner. Typically, those are addressed around architecture or around security considerations with the root server system.

We now believe that most of the rest of the meetings that we should be holding in the future will be open. So, we, by default, going to make all of our meetings open, unless we're meeting with another group that wants it closed or we find that there are security or architectural decisions that need to be held in a closed room. My only slide, right? No, more. Good.

So, that's on the RSSAC policy side within ICANN. The root server operators themselves have also tried to be transparent wherever possible, again, unless security or architectural situations require otherwise.

So, the agendas of all the root server operator meetings, we do meet on a regular basis and the agendas of those have been published for the last few years. I don't remember the start date of those.

There's a public website that contains a list of all of our news and recent events that have occurred on both in terms of service changes in any way or outage reports by certain operators or global statements that we want to make about system-wide events.

We have collaborative reports of major events that occur within that news, as I just said. All of the RSSAC 002 statistics, which are quite extensive, are collected and published by all root server operators and we also publish information within DNS OARC, which is the day in the life of the Internet collection that occurs at least once a year. Lately, because of the KSK roll, it's happened more than once a year, where we publish all the data that's received at the root servers frequently anonymized that can be analyzed by research and data analysts.

Most of the RSOs – I’m not sure about the word all in that sentence, I’m sorry – have webpages. So, we have individual webpages that may have more information about each of the root server operators and the service that they provide to the Internet.

Then, RSSAC can respond to technical questions about the RSS. A couple of years ago, we created an e-mail address called ask-rssac@icann.org where technical questions can be sent regarding the root server system and a response will be provided, and if needed, questions will be sent out to the individual RSOs in order to answer any questions that Internet users might have. Next.

So, most importantly, we want to know is there anything we’re missing? Are you aware of all of these transparency items? I think one of the biggest things that RSSAC has struggled with is trying to make everybody aware of the ways that we publish this much data and are available for outreach and we want interaction with the ICANN community, the technical community, and the Internet at large.

So, if you haven’t been aware of these, we’d love feedback on where we could make sure that this information is getting out, that you’re not seeing it right now.

And what's missing from this list? Is there any ways that you think we're not transparent that we should be? We would certainly love feedback in that way. So, any way that you might help us improve transparency, please come talk to us afterwards or at the microphone at a Q&A session. We'd love to hear about it.

BRAD VERD:

Thank you, Wes. Real quickly, I just want to add that clearly this is an RSSAC meeting, but we are constantly struggling with the confusion between RSSAC and Root Ops, which is why you see that we all represent different root server organizations in our capacity here. It says RSSAC, but we are kind of the front door to the root operators, which is why you see the information that we share here on both sides of the fence between root server operators and RSSAC. I just wanted to point that out. Next slide.

I think that draws our presentation to an end. This is our time for questions. Obviously, this is all of our information that's available on the addresses below. If you're interested in joining the caucus, there's a caucus webpage. Please look into that. Do we have any questions?

WES HARDAKER: If you raise your hand, there's a microphone that we have that we'll come around and find you if anybody has questions.

BRAD VERD: Alright. With no questions, we'll give you some time back. Thank you very much. We are adjourned.

[END OF TRANSCRIPTION]