COPENHAGEN – RSSAC Public Session
Tuesday, March 14, 2017 – 15:15 to 16:45 CET
ICANN58 | Copenhagen, Denmark

TRIPTI SINHA: All right, thank you. Thank you very much. Welcome. This is an update from the RSSAC and the RSSAC stands for the Root Server System Advisory Committee, Committee Advisory to the Board.

Real quick before we get started, I would like to introduce the RSSAC members who are here. I'm Tripti Sinha, co-Chair of RSSAC and I'd just like all the RSSAC members to go around the table and introduce themselves, please.

BRAD VERD: Brad Verd, co-Chair of RSSAC.

KAVEH RANJBAR: I'm Kaveh Ranjbar representing K-Root on RSSAC and RSSAC liaison to the Board.

DANIEL MIGAULT: Daniel Migault. I'm working at Ericsson and I am the liaison of IAB to RSSAC.

TERRY MANDERSON: Terry Manderson, ICANN Root Server Operations.

SUZANNE WOOLF: Suzanne Woolf, Ultimate Root Server Contact for USC.

FRED BAKER: Fred Baker, AFRINIC.

JOHN CRAIN: John Crain with ICANN.

LARS-JOHAN LIMAN: Lars-Johan Liman, Head of Root Server Operations at Netnod.

DUANE WESSELS: Duane Wessels from Verisign. I'm the Root Zone Maintainer Liaison to RSSAC.

WES HARDAKER: Wes Hardaker, University of Southern California.

TRIPTI SINHA: All right, thank you. Carlos, next slide, please?

All right, we're going to go through this with four breakdowns. I'll give you a brief overview of what RSSAC is all about then I'll turn it over to my colleagues and Brad will take you through our publications since ICANN57. And then an update on our current RSSAC and Caucus work and you'll learn about what the Caucus is when I'm speaking, and then we'd like to have some community interaction with you.

Next slide. The clicker's not working, right? Carlos, next slide.

All right, what is the RSSAC? It stands for the Root Server System Advisory Committee, the acronym does, and we are essentially an advisory body to ICANN, the Board and ICANN community, the general Internet community on the operation, administration, security and integrity of the Internet's root server system. It is an extremely well-defined scope and that is exactly what we do, simply advisory in nature on those matters.

This is about the composition of the RSSAC. We have appointed representatives from all the 12 root server operators and each one of us has an alternate as well and we also have liaisons that directly impact our work.

And then we have what's called the RSSAC Caucus. Many of you might be members of the Caucus. Essentially this is a body of

individuals skilled in subject matter and they join the group to do work and they're also appointed by the RSSAC.

Next slide please?

Today the Caucus is made of 85 experts. The way you become a member is they submit a Statement of Interest and you also receive a credit for the work that they do. So we have a call for work and individuals volunteer and the end product is a publication and you are given credit for your work.

The pool of experts, they have expertise in DNS and security in other matters. There's transparency in who does what and there's a framework for getting the work done. We work towards a result. We have deadlines and leaders assigned to the work.

If you are interested, we highly encourage you to apply. Send an e-mail to rssac-membership@icann.org and you would need to submit a Statement of Interest and your background.

Next slide, please.

The next upcoming Caucus meeting is at IETF98 on Sunday, March 26th from 3:30 to 5:00 p.m. and there's something that's cut off. I don't know what that is.

BRAD VERD: That's just the UTC time.

TRIPTI SINHA: Okay, that's UTC time. Just to let you know how often we meet, the Caucus decided that they would like to meet at every even numbered IETF meeting and ICANN AGM meeting. So the next time after Chicago we will meet at the ICANN AGM meeting which, I believe, is in Abu Dhabi.

Next slide, please.

These are just some quick administrative updates. We've got numerous publications and these publications all have numbers associated with it and we have a variety of publications. Some of them are just comments, advisories, reports and procedures and you can actually search against those different types. There's the URL for where our documents live. This is just a list of some of our publications since October of last year.

Go ahead.

BRAD VERD: Just real quick to add that this was a pretty major change recently. We used to have documents by name and it was creating lots of confusion when you were trying to reference something so recently we've gone to a numbering scheme which is what this is here to represent.

TRIPTI SINHA: All right, next slide. And some updates on our membership, Suzanne Woolf used to be our liaison to the Board and she termed out and Kaveh Ranjbar – where's Kaveh? Are you here, Kaveh? He just stepped out. Well, the individual who just stepped out is our representative to the Board and also Brian Reid who represented ISC on the Board and now the root operator is now represented by Fred. Fred, if you could put up your hand. Okay, and that is Fred. Welcome.

Next slide. I will now turn it over to Brad.

BRAD VERD: Hello, everyone? I'm going to run through the recent publications and ongoing work that is happening within RSSAC and then hopefully get to a point where we can have some dialog back and forth.

We have four recent publications: History of the Root Server System, Key Technical Elements, we have a Workshop Report and a Lexicon and again they're all numbered here, RSSAC 23, 24, 25, and 26. I'll go into detail on each shortly, a little bit of detail. Can we go to the next slide, please?

So it's the History document. This was as we were beginning our work and after the reorganization of RSSAC I think in 2014 or

something. It was clear to us that there was different interpretations of how we had gotten here so we thought I would do a rather large effort and try document the history of the root. That's what this document is.

It's a fascinating read and this document actually is the catalyst for the current document that ICANN is undergoing right now, which I think maybe you heard Steve Crocker talk about from stage saying that they're creating a history document of ICANN. This document was a catalyst for that. Once Steve read this, he was really enamored with it and enjoyed it.

It contains, again, a chronological history of the root server system, how it grew, why it grew in the locations it grew in and how the system moved internationally, who it moved internationally to and it describes all those steps that we went through to get there.

Next slide, please.

All right, Key Technical Elements of a Potential Root Operator. Through our different workshop reports, one of the takeaways has been that—one of the things identified and takeaways, obviously, identified that there's no process right now for us adding or removing a root server from the root zone or the [hints] file. We are working towards that process and in order to do that there's a number of things, a number of building blocks

or a foundation that we have to put in place before we can get to that actual point.

One here is Key Technical Elements of a Potential Root Operator. This document expanded on RSSAC001 and RFC 7720. RSSAC001 is expectations of a root server operator. It's a document that explains what's expected, how they're supposed to operate under normal conditions and so forth, whereas the RFC 7720 is very technical and talks about the protocol level of what a root server is supposed to provide.

This document expanded and basically took into a number of different elements including the design, the experience, the networking diversity, and so forth of things that should be looked at when a new root operator is being evaluated. And this was again one of those building blocks so as many building blocks that we're going to be putting in place to get to that overreaching question of how do you add in or remove.

Next slide.

The Workshop Report. RSSAC right now currently we've been working on a cadence of about two workshops a year. We do a spring workshop and a fall workshop. This one was our fall workshop in October that we held at the University of Maryland, where we developed a 50,000 foot mindmap.

When I say mindmap, basically it's changing as we learn more and evaluate things more but this was kind of a line in the sand that showed how we think of a possible model going forward, and it consisted of everything you see here, of empowerment, finance, again designating and removing and then audit and accountability measures. And then obviously out of that came the technical piece, the technical elements of a potential root server.

Other topics that came out of that workshop that have led to work that's underway or has been just recently been completed, the Lexicon, which I'll talk about here in a minute, and then RSSAC and Root Server System Transparency.

It was clear that we needed to look at our transparency level to see where we could improve and in the effort of doing that we actually identified that both RSSAC and the root server system were actually, as it turns out, very transparent in many different ways and we will talk about that here shortly.

Next slide, please.

Our workshop—this is continuing on—we identified that we needed a lexicon. My jump ahead is some of the current work but we just recently finished a lexicon. This was going to be kind of a living document as we start to work through it but we quickly figured out along the way that the community internally

in conversations with the Board, conversations with each other and conversations with everyone, that we weren't working all from the same definition of terms. So we thought that it was very important to sit down and define what these terms were so that we could all start having the same vocabulary, understand what that meaning is. So that came out of the workshop and I think I'll touch on that again here in a minute.

And then again everybody agreed that the designation/ removal function is a necessary process that has to be worked out, but before you can do that there's a number of building blocks that have to be put in place.

Out of that there was consensus in the group to revise and strengthen the RSSAC001 that I mentioned just a moment ago, the expectations of a root server, strengthen, put some stronger terminology, some teeth in the document. Now going back, everyone agreed that there was also a need for an accountability function. And then again I talked about transparency.

Next slide.

The lexicon, again what I just touched on, this just was approved and published. It should be on the website tomorrow or today or... it's there now? Great. This defines a number of very key terms that had different meanings depending on who was

talking about them and what's the context that they were talking about them.

We wanted to resolve that and then it also deprecates a number of terms that you can see here that were just used in a way that just wasn't current so that's why they were deprecated and retired. Again that document was just published. It's on the RSSAC webpage so feel free to check that out.

Next slide.

All right, current work, let's jump into this. These are ongoing work parties that are happening today.

Next slide, please. This is the Root Server Naming Scheme. This work party has been in existence for a little bit. They've been making some really interesting progress. It's ongoing right now.

This one was a result of some of the work we did on the history document as we were identifying what the history was. One of the questions that came up was the naming scheme used for the root servers, a.root-servets.net, b.root-servers.net through m, that naming scheme was put in a place almost 20 years ago and so the question was obviously it still works today but is it the best thing to be done today? So we wanted to ask these questions. Is there a better way to do it? Is there a more efficient way? Do we just need to rethink this? And so that's what the

group got together and gone out. They've been doing a number of tests and evaluations and they're supposed to come back with a recommendation. Given a blank slate, starting over again, what is the recommendation on how it would be?

Key part of that document, though, is that we should talk about what the risk analysis is of changing as it sits today. Are there new risks we would be introducing which is something that we want to avoid and if so, what are they and we should talk through them.

Next slide.

Another work party, Distribution of Anycast Instances, I'm actually going to turn this over to the work party leader, if I may, and he can give you a quick update. Kaveh, who was not here earlier, this is our liaison to the Board. He is our work party leader for this effort so I'll give the mic to him.

KAVEH RANJBAR: Thank you very much, Brad. I'm so sorry for not being here. I had a phone call to attend to. Basically this work party here, the first one, if you're interested in what I'm going to explain, please drop a line to me or talk to any one of us and we would be more than happy to have you on board because there is a lot of work.

This work is basically that the title reads Best Practices of Distribution of Anycast Instances of the Root Name Service, but there are actually four questions in this work party. One is we wanted to find a measure—I'm paraphrasing what's written there—to be able to see if root server operators are doing a good job. That's basically the gist of the first question. And latency is the main measure we wanted to focus on but there is a group of people and actually one of the discussions is should we look into latency or not. So the way the work party is working is not limited to the question. They're also questioning the actual question and trying to find a real answer.

Another one is adding more instances in more topologically diverse locations to make the system more resilient to DDoS or not, which I think is clear enough. All of us have a feeling but we never had any empirical evidence to basically support that so we are trying to see if we can actually find proof that it adds to resilience or not.

The third one which I think is also a very important question is should root operators basically coordinate location of anycast instances or not? Because until now the system, one way or another, has grown more naturally but maybe it's better to actually have a coordination and say, "Hey, I have few nodes in South Africa so maybe you should focus on somewhere else." We

haven't done that at least systematically before but we want to know if we should do that or not.

Finally, is there any technological risk that only a subset of operators versus all or majority of root operators deploy anycast instances – will this cause any risk if everyone is anycast or not or only partially anycast?

The reason that I encourage you to please join the work party is first of all there is a lot of work but second, from the looks of it and especially, for example, the question three which is much more visible, if root operators have to coordinate or not, the answers to this questions will be fundamental to the future work of RSSAC, which you will also hear about. Because there is a lot of other stuff which we have to build on top of this documents and this answers to know where we are going with the future of root server system which is basically what RSSAC is focusing on right now. Thank you.

BRAD VERD: Great, thank you Kaveh. So again that's an ongoing work party right now.

Next slide, please.

This one we're just in the very beginning phases of creating this work party so a call will be going out to the Caucus very shortly,

and I'm going to turn this over to the gentleman who is our work party leader on this, Lars Liman.

LARS-JOHAN LIMAN: Thank you. Yes, this is a rather small effort, I hope. You never know until you start to drill down into it. The root server operators collect statistics for various purposes. Sometimes we share the statistics with others. Typically we do a one-time-per-year collection for 48 hours where we collect all the incoming queries to all the nodes on the root servers and we upload them to a central repository at the DNS that is the Operations Analysis and Research Center, DNS OARC, who has this giant databank. There it can be used by researchers who want to find trends and do DNS analysis on these queries.

There is a slight problem with that and that is that the packets include the source IP address of the entity that sent the query to the root name server and that is a bit of an integrity problem.

Some root server operators have found themselves in the situation that they don't want to share the IP address of the entity that sent the query so to circumvent the problem, until now we've been anonymalizing the queries so we randomize the address and changed the IP address before we upload the data to the [repository].

That has to be done in a consistent way because one of the things the research may want to look at is if there are multiple queries coming from the same entity and what the queries contain. There is a path there where you can look for and find problems with the resolver software so one must be able to identify the same source but not specifically which source. But it repeats the queries is the significant part.

There are multiple ways to do this anonymization and this work party which hasn't been fired up yet, the intent is to look at whether harmonizing these procedures so that all the root server operators who do anonymization do it in the same way, that they use the same mathematical rules so that you can find a query source in the data from different root server operators and recognize that this is the same source because that's not obviously the case today.

So that's the first step. And if they agree that this is actually something that we want to do to harmonize, we will also like them to suggest a mathematical method to do that so that we know what to use and also to consider whether to recommend that all root server operators do this, not only the ones that are under legal obligation to anonymize or have some other concern about this.

So today it varies in the repository. Some of the data is anonymized, some is not and we want this work party to investigate whether we should recommend that all do the same type of anonymization so that it can be compared all over the line. So that's basically it.

BRAD VERD: Thank you, Liman. There always seems to be interest in retailer and right along those—I'm sorry, next slide—right long those lines there always seems to be development or effort or need for tools to consume and use that data, and right now we're starting an effort that Wes is going to describe. Wes Hardaker from [ISCOS].

WES HARDAKER: Thanks. ISI.

BRAD VERD: Apologies.

WES HARDAKER: No worries. One of the things now that most of the root operators are publishing data with respect to operations on a day-to-day basis on statistics regarding it, they're all defined in

our seg 002 which is one of the earlier documents in the RSSAC history.

We realized that it would be helpful if researchers had the ability to coordinate with code and to promote code reuse some things like that. As these results in particular are likely to be a foundation for a lot of future research, it'd be much better if we could have a common base that everybody can work from.

In that light, we have started a project within the RSSAC Caucus. Again anybody can be a member of the Caucus if you want to contribute. We'd love your help too and our goal is to produce that common code base for at least [short] coding resources regardless of the others, actually RSSAC002 related research or other research into the root server system that people are writing code for.

There is a hub repository that we've created. There's nothing in it right now. We will be discussing the format and the layout and how we're going to make use of that resource in Chicago in a couple of weeks so after that time we hope that stuff will start appearing there. And we have a mailing list as well that members will contribute to and have a discussion about. And then if anybody wants to join the RSSAC tools list, it is a Caucus oriented activity but you can contact me at that address at hardaker@isi.edu if you're interested in participating. Thanks.

BRAD VERD: If I may, Chicago in a couple of weeks is again the Caucus meeting that is occurring in Chicago at the IUTF.

WES HARDAKER: Yes, thank you. One quick addition, you don't have to be at the Chicago meeting to help out. That's where we'll talk about some organization and some initial ideas and the projects that people want to do. If you're not going to attend the Chicago meeting, that's fine. Feel free to send me details about stuff you want to do and how you want to contribute and code you might want to work on and I'll make sure that that gets discussed in Chicago as well.

BRAD VERD: Remote participation will be available for the Chicago meeting. All right, next slide please. I think we're going to go right back to Wes for community interaction.

SUZANNE WOOLF: Wes, before you proceed, I was just informed that we have three of our colleagues who have joined remotely so I'd like to acknowledge them. Online we have Jerry [inaudible] from the University of Maryland, Kevin Jones from NASA, and Kevin

Wright from the Department of Defense in the U.S. So welcome to them. Thank you.

WES HARDAKER: One of the goals that RSSAC and the root operators have had is to increase our transparency and to publish more and more of what we do and how we do it. And to that end, this is the list of, I guess, action items that we have done in the past number of years and so I'll run through them really briefly.

On the RSSAC side of things, we've established the Caucus which allows everybody to participate in the ICANN half of the policy making actions within RSSAC, as well as we publish all of our minutes and workshop reports.

We have a public RSSAC in Caucus calendars where everybody can see what's happening and when. The RSSAC public meetings are like this one where we interact with and we want comments and we solicit feedback.

We have meetings with other ICANN community groups. We're constantly trying to reach out and make sure that we're in touch with all of the other communities within ICANN.

We hold tutorials. What is a root server tutorial is held twice at every ICANN meeting. If you missed it this time, please come to the next one. It's at the beginning of the next ICANN meeting. We

have liaison relationships with a lot of other external bodies and our operational procedures are publicly available and documented in our RSSAC000 document.

On the root server operations side of things, any meetings that are held with them, published meetings are available for reading. As I mentioned a few minutes ago, the RSSAC002 statistics play a major role in showing the number of queries and all sorts of statistical information about how the root server system operates.

The root server operators all participate in RSSAC by delegate and we have a public webpage which is www.rootservers.org and on that page you can find a lot of information about the root server system as a whole where its 600 instances lie, as well as links to individual webpages for those root server operators that have individual pages about their operation as well.

There are some operators that have public letters that say IANA on the slide because you can't read it because it's covered up by a white blob. It's actually ICANN so the MoUs that exists directly with ICANN IANA.

And then we produce collaborative reports on major events so if there is a major Internet event that affects the root server system, there's a common report that triages what happened and what actions were taken. And then finally RSSAC has agreed

to be the front end to any questions about the root server system so if anybody has any technical questions, if you don't want to bring it up at the microphone today, you're welcome to write ask-rssac@icann.org and those technical questions will be answered and delegated out to root server operators who may need to help answer it but we'll try to make sure all questions are properly answered. Thank you.

BRAD VERD: Thank you, Wes. Just a clarification because I know it gets confusing for me also and I imagine it probably happens for others, the RSSAC members wear many different hats including root server operators and RSSAC so when we heard Wes talk about we in both columns, it's with different hats on, just to make that very clear.

RSSAC again, as we stated in the scope earlier, we give advice to the Board about the root server system. Obviously as a root server operator, you have a different world and RSSAC doesn't play in that world.

WES HARDAKER: Thank you very much, Brad. That's absolutely true.

BRAD VERD: Great. Next slide, please. Again hopefully this'll lead into our dialog. This is a result of the last slide that you just saw that showed the interactions with the community from both RSSAC and the root server operators.

A couple of questions: Were you aware of all these things and really what's missing and what further can we do, is what we're asking the community. If you can think of anything, please let us know either individually to the mailing list or come up to the mic.

Next slide, please.

And that comes to the end. Here is obviously our main webpage where our publications are. Again, if you're interested in joining the Caucus, there's a link there that describes the Caucus and the membership e-mail address. With that, I'd like to turn it over to you with any questions that you may have for us that we may try to answer.

Carlos, are there any questions online?

CARLOS REYES: No questions remotely.

BRAD VERD: If there's no questions, I will give you back some time but we would encourage you to talk to us one-on-one, grab us as a group, whatever works, whatever you feel comfortable with. If there are no questions I will close this meeting and adjourn it. Thank you. Have a wonderful day.

[END OF TRANSCRIPTION]