
BARCELONA – RSSAC Information Session
Tuesday, October 23, 2018 – 09:00 to 10:15 CEST
ICANN63 | Barcelona, Spain

BRAD VERD:

We're working through some technical issues. We'll be starting here in just a minute. It'd be great if anybody could – yeah, there's plenty of room at the table, if you guys can join us, so I'm not speaking behind me. Otherwise, I'm happy to stand up with a mic and walk around and talk so you're not listening to my back. Thank you all. My spine specialist is very happy that I don't have to spin around and try to ...

Again, we're working on the phone bridge for the Adobe Connect, so just give us one second.

All right. Thank you, all, for being here. I'm Brad Verd, Co-Chair of RSSAC. Welcome to the RSSAC public session here at ICANN 63. We're going to run through our activities update. I'll give you – let's see. We're going to cover these five things. We'll do a quick RSSAC overview. We'll run through the publications since ICANN 62. We will go over the proposed governance model, which was RSSAC037, which was submitted to the Board prior to ICANN 62. But, because it's so substantive, we thought we'd cover it again here. Then, we'll cover current work and community interaction.

So, the overview. This will go quick. What is RSSAC? So, this is taken right out of our bylaw, basically, which is, "The role of the Root Server System Advisory Committee is to advise the ICANN community and the

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.

Board on matters related to the operation, administration, security, and integrity of the Internet’s root system.”

So, what does that mean? Well, it’s just a very narrow scope. We are not a policy body. We are not an operations body. We are a body appointed by the Board to advise the Board on the root server system.

So, with that, I’ll give you the background here. So, RSSAC-appointed representatives from the twelve root server operators. Each of the twelve root server operators or organizations, in addition to the appointed representative, has an alternate. In addition to the twelve appointed representatives and the alternatives, we have liaisons to a number of community bodies. Those include: we have liaisons to SSAC. We have a liaison to the Board. We have a liaison to the IAB. We have a liaison to the Root Zone Maintainer ... I’m missing –

UNIDENTIFIED MALE: PTI.

BRAD VERD: Oh, I’m sorry. We have a liaison to CSC, a liaison to RZERC. There’s just a number of them. You can see that this is how we make sure that we’re involved and engaged in the right parts of the community that need our support and expertise.

In addition, there’s the RSSAC Caucus, which is a body of volunteer subject matter experts appointed by RSSAC. The next slide here is – right now, the Caucus is in excess of 100 DNS experts. All of them have

provided SOIs. All work that comes out of the Caucus is credited to the individuals who work on it. So, if you are a member of the Caucus and contribute work on any of our statements or our documents, you get full credit for it. This is not something that is handed up to anybody else.

What's the purpose of the Caucus? Obviously, it's a pool of experts who try to bring all that expertise together so that any advice that RSSAC is giving to the board is fully-encompassing of what it should be.

Transparency, as I stated early, of who does the work. Again, if you want your name on a document, it's very easy to do here in RSSAC. It's just: be engaged and you will get full credit.

We have a framework for getting things done. We do statements of work. They're sent out for a call for a work party. A work party is formed. We identify a work party leader. The leader runs the calls. Then, in addition, we have, from RSSAC, a work party shepherd, just to make sure that communication is happening. Obviously, we put together a plan for any work party work, and with deadlines, and we monitor it through.

If you want to be a member or you're interested in looking more into it, there's an e-mail address there, rssac-membership@icann.org. Please don't hesitate to reach out.

The Caucus, by their own decision – so, we have talked to them extensively about when they want to meet, and they have chosen to meet at every AGM. So, we have a Caucus meeting following this

meeting, not in this room – a different room. But, there’s a Caucus meeting right after this.

They’ve also chosen to meet at every even-numbered IETF, since a lot of the technical experts that are in the Caucus also attend the IETF meetings. It’s getting to a critical mass of people in one place and taking advantage of that. So, we do every other IETF, and our next one will be in March in Prague.

So, that’s the Caucus in RSSAC.

Let me run through ... is this you? I’m sorry. I’m going to hand this over to my colleague, Liman, who’s going to run through the publications, and it will come back to me shortly.

LARS JOHAN-LIMAN:

Thank you. My name is Lars Johan-Liman. I work for Netnod and [I am] one of the members of RSSAC. Since ICANN 62, the last ICANN meeting, we’ve published three documents in our RSSAC documents areas, one a statement regarding the KSK rollover plan, which was published before the rollover, and one regarding recommendations on the anonymization processes for source IP addresses when we submit them for analysis, and one which is reflection over the recent organization review that we underwent. All go through all three of these.

So, the first one was a statement in response to a request from the Board, where the Board reached out to a number of the communities and advisory committees and asked for advice regarding the planned

KSK rollover for the DNSSEC key signing key for the root zone. Obviously, the root server operators and the RSSAC were one key community to reach out to.

So, we looked at it and we found two items where we recommended that the Board just do a last check to see that everything was in line with what they planned for. One was potentially increasing traffic to the root server system, and one was to review the published recovery plan. So, should something go wrong, the rollback plan was actually well-oiled and thought through properly.

But, in conclusion, we said that RSSAC did not see a technical reason for these items to cause any delay in resuming the KSK rollover plan, according to the plan schedules. And, that eventually happened: the KSK was rolled on October 11th. As far as we know, everything went well, except for the small, odd thing here and there, which everyone expected. But, no substantial problems that I'm aware of, at least.

The next document was actually a planned work party document. So, this was one of the issues that we asked for help from the Caucus. So, regularly, the root server operators collect incoming DNS queries to the root servers and, at least on a yearly basis, we report all these recorded queries to a central repository every year. This is called A Day in the Life of the Internet, or the DITL collection. They're stored there for future reference so that a researcher can look at long-term trends in DNS traffics. Not only do root server operators do this, but it's voluntary thing, so other DNS operators do it as well; for instance, TLD operators and also one or two resolver operators [of Iraq].

The thing is that we report the queries to this central repository, and the queries contain the source IP address of the client. Some operators are uncomfortable or even prohibited by law to provide this detailed information about the client.

So, a few of us, Netnod included, have anonymized the source IP address. This was a request to the Caucus to investigate whether there was reason to select a specific anonymization process that would be better than the others and also whether to harmonize that between the root server operators so that we would do it in the same way.

This document provides a bit of guidance regarding algorithms and such for anonymization and discusses their pros and cons and so on.

So, the three recommendations that come out of this document is that root server operators should consider the pros and cons of harmonization and anonymization, and also that they should do that individually. So, it's up to each operator to make these decisions.

But, they also recommended that the autonomous system number, the identifier that identifies the Internet service provider that is closest to the end client, that AS number, also be part of what's submitted, as long as that doesn't identify the client too much. If you have a large ISP, that information is not very sensitive, but if you run your own network and have an autonomous system number, it can identify you fairly closely. So, this is something we will have to look into.

So, that's actually somewhat difficult. It would change the procedures for how we submit data today, but it's not impossible. So, that's something that the root server operators will have to look into.

The last document is a consequence of the periodic organizational review that RSSAC underwent. All the constituencies and parts of ICANN undergo regular organizational reviews, where a review team evaluates how the sub-organization – in this case, RSSAC – function as a committee. Does it work well? Does it follow its procedures? Does it follow its bylaws? Blah, blah, blah. This happens on a regular basis for all parts of ICANN, and now it was our turn.

We issued a document with reflections on how this review was conducted. We felt that some of the things that were evaluated were kind of possibly a bit off topic, so we've made a few recommendations for future organizational reviews in the ICANN community. Those are, in short, to make sure that a review addresses the responsibilities of the ICANN organization for the organizational review process so that it's very clear what the review team is actually going to assess.

The document also provides advice on writing the request for proposals and selecting the reviewer because the reviewer is supposed to be an independent examiner, which is something that is procured from outside ICANN to have an unbiased view of what's going on.

The document also includes guiding principles for how an organizational review should be conducted. We felt that, as I said – maybe looked at some things that didn't pertain to an organizational review.

It also provides five recommendations to ICANN regarding these organizational reviews and how to conduct them. These are the five recommendations.

So, we think that the ICANN organization should define organizational reviews in a more strict wording so that it's clear what's about to happen and that the old parties know what's going to happen, and to do that thing in a fairly crisp and tight order.

It should also document the intent of the organizational review, what information it hopes to obtain and how that information will be used down the line.

It should continue to use its RFP process that it has, but it should assure that the independent examiner, the IE, are experts in assessment frameworks and methodologies and that they are not from the ICANN community. So, even if someone operates a company that may seem suitable, but where the individuals are involved in the ICANN community, in the ICANN work, maybe that's not the best ones to select.

Yes, there should be a continuous connection between the independent examiner and the ICANN organization with checkpoints and so on so that the guidelines above are carefully followed.

And, at the conclusion, when an organizational review is finished, the ICANN org should report on the process, how it has transpired, and, if there are any lessons learned from this particular review, [they should also be included] as feedback to future reviews.

So, those are the three documents we've released since the last ICANN meeting. I think it's back to you, right?

BRAD VERD: Actually, Ryan Stephenson is going to talk about the governance model. So, go ahead, Ryan.

RYAN STEPHENSON: Good morning. I'm Ryan Stephenson. I'm with the Department of Defense.

BRAD VERD: Ryan, just give me the queue and I'll do the slides for you.

RYAN STEPHENSON: Sure. No problem. Thank you. So, for the past three years, RSSAC has been working on a proposed governance model for the DNS root server system.

In the past, the Internet has grown, so it's thousands of times larger than what the root server system for DNS was originally created for. There's billions of posts, billions of users, new governance structures, and new business models. All this places new expectations on the infrastructure. So, we can all agree that the Internet is pretty important.

So, throughout it all, the RSS, the Root Server System, has largely maintained the same organizational structures through all of these changes. It has adapted to growth and provides a resilient service.

But, now, it's kind of time for the root server system to adopt a new governance model and structures to evolve to the business models that are out there for today. So, this is to meet accountability and transparency.

So, what we've done in this proposal for a governance model is we defined eleven principles for the root server system and also the root server operators. We demonstrated, also, how the governance model will work with removal for poor performance, voluntary removal, adding an operator, catastrophic shutdown, and also for a rogue operator.

Next slide, please. So, if you've been attending any of the RSSAC sessions, you probably might be familiar with this slide. Here, this slide shows a TLD [operator]. [If] they request any changes, they'll go to IANA, which will batch out the changes and send them over to the root zone maintainer, which will distribute the root zone file to the operators. Then, the operators publish out the zone file, where DNS resolvers then go to each operator's consolation.

So, where the model lies within this is within the resolution portion. This is a kind of tightly scoped and it's well-defined for where this model resides.

So, go ahead to the next slide, please. Right here we have the eleven principles. So, these eleven principles guide the development and operation of the root server system and the root server operators. RSSAC uses these principles and believes that they should remain at the

core of the governance model. So, we'll quickly go over some of these principles here.

Some of the principles – like, for example: to remain a global network, the Internet requires unique name space. IANA is recognized as the source of the DNS root data. The root server system must remain a stable, reliable, and resilient platform for the DNS service for all users.

Diversity of the root server operations is a strength of the overall system. What that means is that we have different organizational structures, we have different funding models, and we have different technologies and topologies and so forth. This diversity creates the strength of the root server system in general.

Architectural changes should result from the technical evolution, demonstrated by technical need. In other words, let technical aspects run the root server system.

It recognizes that the IETF defines the operation of the DNS protocol.

The other principles? RSOs must operate with an integrity, an ethos, demonstrating commitment to the common good of the overall Internet. We must be transparent, and RSOs are working to strive to maintain that. We collaborate with the community at ICANNs, at IETFs, and the stakeholder community.

One of the aspects of the diversity is that the RSOs remain autonomous and independent. What that means is that we kind of follow our own paths for the good of the Internet so as to not, so to speak, copy another root server operator.

And, RSOs must remain neutral and impartial, that being that we're apolitical. We just serve the root zone file in its entirety with no manipulations or anything.

Next slide, please. So, with the proposed governance models here, the model consists of, like, three cores: the governance, the DNS operations, and the onboarding and offboarding. The model was built on separation of functions, avoidance of conflicts of interests, and transparency and accountability.

So, within a model here, we identified who the stakeholders are. It's at the top of the box there. You can see that it's the ICANN community, the IETF IAB, and the root server operators.

Within a model, we identified several functions, one being the strategic architectural and policy function. So, this function is the structure ... sorry. I have my notes here, too ... that offers guidance concerning the root server system.

We also have the financial function. The financial function provides finances for research and development, operations, and contingency.

We also have the designation removal – oh. Let me actually go to the performance monitoring and measurement function. This is where it will take input from the strategic architectural and policy function and measure up the root server operators and also monitor the root server system as a whole.

Then, we have the designation removal function. This is the function that would make the determination as to who gets removed and who gets added into the root zone files.

We also have a Secretariat function. This is mainly used for the root server operators. This function kind of acts as, really, a secretary, so to speak. It assists the root sever operators with their meetings and any tools that the root server operators may use.

Within this, there's performance metrics that the root server operators must maintain to be added into the DNS root sources, which are the three files there: the root zone, the root hints file, and root-servers.net zone. This is where an operator would then be removed or then be designated.

Next slide, please. So, with this, we also have RSSAC038. RSSAC038 accompanies RSSAC037, and it's to initiate a process to produce a final version of the model based on RSSAC037.

Within RSSAC037, we identified a way to measure the cost of the root server system by using something called BPQ (Bandwidth/Packets per second/Queries per second).

So, here we also estimated the cost of the root server system in developing the model. So, the model was presented to the Board prior to ICANN 62, and we're waiting for feedback from the Board as to the execution of the model. And, of course, in RSSAC038, the final version of the model should be based upon the principles that were laid out within RSSAC037.

If anybody has anything else to add from the RSSAC, please do so.

BRAD VERD:

Thank you, Ryan. I think I'll just add, again, that this was a substantial effort put forth by RSSAC. It was submitted to the Board, and we are waiting on the Board response right now as to what was get back here. This was intended from RSSAC to be the beginning of a conversation with the community. So, we expect feedback. We want feedback. So, we look forward to hearing from the Board and from everyone else.

All right. Moving on, this is the update on the Caucus work? ... current work and Caucus work. I apologize. So, we are currently working on completing our second organizational review. I know that Liman talked about a statement that we published with five recommendations regarding our experience with the review process. The review continues.

So, obviously, in June, the examiner finished its work. Earlier this month, we, RSSAC delivered our feasibility assessment and initial implementation plan. It's a bit of a tongue-twister. I did not name it.

We gave that to the OEC, and we are waiting on feedback from them. So, they're now considering the final report.

The next step there is we sit down with the independent examiner and the Board, come up with a final report, which will then be published as an RSSAC document, and we work through with the Board any recommendations that we need to implement. So, that is still underway.

Secondly, we have two work parties underway with the Caucus right now. The first one here is the service coverage of the root server system. We are trying to explore the concept of accessibility to the root server system.

Actually, Liman, did you want to talk to this one?

LARS JOHAN-LIMAN: [inaudible]

BRAD VERD: I'm sorry ...

LARS JOHAN-LIMAN: Never mind. So, I'm the shepherd for this work. This is work that has just begun. We had our first telephone conference just two weeks ago, I think. So, this is work where the root server operators receive comments regarding service coverage, saying that there are areas of the Internet that don't receive, so to speak, proper service from the root server operators.

Root server operators are always keen to address that, but, in order to do that in a good way, we need to understand what the problems are. So, we have now asked the Caucus to help us investigate – I should have the right slide in front of me – how to determine this, how to get the knowledge, how to assess what's going on on the Internet.

So, in order to do that, we first need to discover and understand what indicators, what technical measures, can we do to try to find where the problems are? Where should we measure, what should we measure, and how should we measure in order to understand how the system is seen and viewed in different parts of the Internet?

When I say “different parts,” I don’t necessarily talk about geographical parts because the Internet doesn’t follow national borders and geography all that much, so it’s a topological view. How is the network connected together, and which parts of this connected network don’t work, and why not?

If you start to drill down into the details, it turns out to be extremely complicated how the Internet works. I wouldn’t say sheer luck, but the fact that it actually works is astonishing.

But, once we know what to measure and how to measure, we also ask this work party to take the next step: to suggest procedures and tools that we could use to actually perform the measurements. And, then, as the third step, to actually perform these measurements. So, this is a bit of undertaking to actually do these measurements and help us find topological areas in the network which don’t receive as good a service from the root server system as a whole that we would like to provide.

Then, as the last step: to recommend to the root server operators and the Internet community how to enhance the service coverage of the root server system.

So, this is actually a lot of rather large work, but we hope that it will lead to the root server operators' having better tools and a better understanding of which parts of the Internet are not served at the level that you could expect or where service could be improved and then also get a better understanding for how and where to improve this.

I'm also for the next one, aren't I?

The next work party that also I believe was launched fairly recently is a work party regarding modern resolver behavior. The resolver is the counterpart, the client side of the communication with the root servers. We would like to investigate how modern recursive DNS resolvers interact with authoritative servers, and especially the root server system.

Again, the root server operators would like to provide as good a service as possible, but if we don't fully understand how the clients behave, it's very difficult to meet that from the server side, from the authoritative side. So, we really want to understand how the resolvers function and how they work in order to adapt the root server system so that it gives the best possible service to these resolver systems.

So, the work party is to potentially recommend changes to the DNS protocol. That's something that we will have to funnel back to the IETF, but also to implementations. That would probably go for both the resolver side and possibly for the server side as well. These are two parties that need to meet. In order for that to function well, there has to be a bit of negotiation.

Fortunately, most of the implementers, all systems with large deployment, are known in the community and are very open to discussion/dialogue regarding these things. But, we need to do this work to find where we need to tune things, and also, potentially, at the far end, create advice to the ICANN Board or other ICANN constituencies to benefit the stability of the DNS on the root server system.

Is it Ryan from here, I think? Tools.

BRAD VERD: It's Wes.

LARS JOHAN-LIMAN: Wes.

WES HARDAKER: Wes Hardaker, USC-ISI. One of the things that we have been trying to do with all the results of RSSAC work regardless of RSSAC Caucus work, regardless of what went into various documents, is to make sure that all of the results are archived on GitHub. So, tools are written, or sometimes data – whatever it might be, it's a fairly open type of thing. But, we have a GitHub repository for the RSSAC Caucus. It's [GitHub.com/rssac-caucus](https://github.com/rssac-caucus). There you'll find, basically, the results of anything that we've done in the past, or even current work that's ongoing. We have six or seven repositories in there now. I suspect that the two new work efforts that Liman just talked about will also end with

repositories there since we're likely to be writing code and writing simulations and things like that.

We want the results of what we do to be reproducible and to be read again in the future. Resolvers change over time, so we want to be able to remeasure stuff in the future after resolvers have been updated, according to the last work, for example.

So, in there, you'll find lots of information. One of the biggest ones in my mind is: all of the RSSAC002 data from all of the various root organizations is archived there. So, rather than have to go fetch it from each particular location, you can actually just download it from GitHub all in one shot. That makes it a whole lot easier to fetch all the current data. Thank you.

BRAD VERD: Thank you, Wes. Thank you, Liman ... oh, Wes. Did you want to continue this?

WES HARDAKER: Ryan was going to do this one, actually.

BRAD VERD: I apologize.

RYAN STEPHENSON: Hey. Good morning again. I'm Ryan Stephenson with the Department of Defense, RSSAC. So, I'm going to go ahead and cover transparency.

This is one of the things that the RSSAC continually strives for: to be transparent.

Starting at ICANN 62 in Panama, we've opened up our meetings to the public. The only that the meetings may not be open to the public is when we're discussing sensitive RSO or Root Server System information. However, we strive, though, to have all of our meetings transparent as much as possible.

So, with transparency, we publish the minutes of our meetings. I'm not going to go over the links, but you see that the links are up there and also in this slide deck.

Our publications. If you go out to the RSSAC web page on ICANN, you'll find all of our publications for meeting minutes, publications, and even the RSSAC000 operational procedures with how the RSSAC functions.

Also, at each ICANN, we offer the root server system tutorial, as well this RSSAC brief, which we're in right now.

The RSSAC server system tutorial is a great session because there it kind of uncovers and demystifies the root server system.

Next slide, please. So, for root server operators, we have a website, root-servers.org, and we publish the agenda of the root server operator meetings to that website. So, if you want to see what the root server operators discuss, we go ahead and publish that to our website.

If there's a major event that affected the root server system, we sort of collaborate on reports and we will publish them to that same website

as well. There's a news section. So, periodically check that, just in case there's an update to one of the root server operators or if there was an event for the root server system.

Each of the root server operators have a public webpage. So, this is a web page, kind of like a homepage, as to updates that the root server operator itself is doing, either to better the root server system or some news events for the organization of the root server operator.

We publish RSSAC002 statistics to the root server operator webpage, as well as on the home pages itself. There's a link – well, actually on root-servers.org – to the homepage or to the page that has these statistics. There's several sets of statistics that are pretty interesting to find out what the root server system or the root server operator is functioning.

Of course, you could go out to DNS-OARC, and there's statistics out there as well, as well as the Day in the Life of the DNS.

If you have any questions about the RSSAC or root server operators specifically, you could reach out to the RSSAC by going to ask-rssac@ICANN.org.

So, these are ways that we're striving to become more transparent, but the list isn't fully complete because we need, actually, input from the community. So, this is where we go back to the community – next slide, please – with, were you aware of these transparency items, and how can we be more transparent? We would really love to hear from the community and work with the community on where we can become more transparent.

That's about it. If there's anything else anyone would like to add for that, please do so.

BRAD VERD:

Thank you, Ryan. I'd like to add just real quickly – let me go back. So, again, this is an RSSAC session. As it turns out, the members of RSSAC encompass the root operators, so we were different hats at different places. This is not a root operators meeting. We don't represent here and share the views of the different root operators.

But, what we're trying to share here is the different transparency for RSSAC to root ops that have happened, that are implemented, that are going on. So, that's why that is shared.

Also, it was at least two years ago, I think, that RSSAC agreed to be the window into root operations because there wasn't any. It was a bit of a black box, a bit of an unknown. There was no way to reach out and contact any of the root operators. So, that's, again, why ask-rssac@ICANN.org was created; so that, if there was a question about a particular root operator or something that happened, it would come into RSSAC, and RSSAC would act as the intermediary and send it off to the root operators. So, that's why that's there.

UNIDENTIFIED MALE:

I noticed you publish the agenda for the meetings, but do you publish the minutes?

BRAD VERD: I'm sorry. For which meeting are we talking about?

UNIDENTIFIED MALE: The root operators meetings.

BRAD VERD: Yeah. Minutes are published.

UNIDENTIFIED MALE: Okay. And, the meetings at the IETF, are they public now as well?

BRAD VERD: I'm sorry?

UNIDENTIFIED MALE: The meetings at the IETF. Are they public meetings?

BRAD VERD: No, they're not.

UNIDENTIFIED MALE: They're still closed.

BRAD VERD: Those are operational meetings.

UNIDENTIFIED MALE: Okay.

UNIDENTIFIED MALE: [Sorry. We don't publish the minutes, do we?]

BRAD VERD: Yeah.

UNIDENTIFIED MALE: Okay.

BRAD VERD: All right. Now, we're at the questions section, so thank you for that. Up here is all of our different contact areas – the webpage, the publications. We have a FAQ. We've created a FAQ. The round of questions that RSSAC gets is rather predictable from, I want to say, new members asking questions about root operations, root operators. The FAQ covers all of those questions that we've been asked with very detailed answers and links. So, please take a look at that.

Then, obviously, there's more information about the Caucus on the Caucus page.

So, with that, we'll open it up to questions.

No questions? If there are no questions, we will adjourn.

All right.

LARS JOHAN-LIMAN: Going once ...

BRAD VERD: Thank you very much. We are adjourned.

[END OF TRANSCRIPTION]