SAN JUAN – RSSAC Session
Tuesday, March 13, 2018 – 15:15 to 16:45 AST
ICANN61 | San Juan, Puerto Rico

TRIPTI SINHA: Welcome, everyone. This is the RSSAC Open Session where you're going to learn all about RSSAC and the Caucus and its related activities. Let's start with introductions. I'll start. We'll go counterclockwise and then into the back and then go and check and see who's online.

So I'm Tripti Sinha, University of Maryland, Co-Chair of RSSAC.

MAURICIO VERGARA: Mauricio Vergara from ICANN.

DARREN KARA: Darren Kara, ICANN and RSSAC Caucus member.

BRAD VERD: Brad Verd, RSSAC Co-Chair.

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

VICTORIA RISK: Vicky Risk, ISC.

FRED BAKER: Fred Baker, RSSAC member.

JEFF OSBORN: Jeff Osborn, RSSAC.

NAELA SARRAS: Naela Sarras, IANA Functions Operator Liaison to RSSAC and also an ICANN staff member.

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

WES HARDAKER: Wes Hardaker, University of Southern California.

MICHAEL CASADEVALL: Michael Casadevall, independent freelancer.

ALEJANDRO ACOSTA: Alejandro Acosta, RSSAC Caucus member and I work for LACNIC.

DANIEL MIGAULT: Daniel Migault, RSSAC IAB Liaison.

ANDREW MCCONACHIE: Andrew McConachie, ICANN staff supporting RSSAC.

MARIO ALEMAN: Mario Aleman, ICANN staff supporting RSSAC.

MING-JEN WU: Ming-Jen Wu, PSWG member.

UNIDENTIFIED MALE: [Inaudible] [for ICANN].

KAZUNORI FUJIWARA: Kazunori Fujiwara from JPRS.

KEITH MITCHELL: Keith Mitchell, DNS-OARC, RSSAC Caucus member.

UNIDENTIFIED MALE: [Inaudible] University, Pakistan.

TOBIAS BOLLIGER: Tobias Bolliger with Federal Police and Public Safety Working Group.

[SOPHIE]: I'm [Sophie] from Taiwan. I'm just an observer.

TRIPTI SINHA: Thank you, and we've got many others online. Welcome to all of you. We'll spend the next 60 to 90 minutes going over what the RSSAC is, what the Caucus is, and the work we've accomplished thus far.

So let's start running the slide deck. Next slide, please. All right, thank you.

So we'll, I've broken the presentation up into four sections, a quick overview of the RSSAC, we'll move on to a discussion of our publication since ICANN60, then updates on current work, and then interaction with our community. Next, please. Next.

So the root server system is an artifact of the board and our mission is to advise the ICANN community and the board on matters relating to the operations, the administration, security and the integrity of the Internet's root server system. And this is an extremely well-defined scope. There should be no confusion what exactly we've been brought together to do and we are

purely advisory. The operative word here is “advisory to the board”. Next slide.

We have 12 appointed representatives from the 12 root server organizations that currently operate the root and we have primaries, 12 primary representatives from these organizations and an alternate each. And in addition to this, we have five liaisons. And these liaisons serve different functions and we’ll hear from them as we continue along this presentation.

We also have what we consider an extension of this, of the RSSAC and it’s a body of subject matter experts. We delve into a fair amount of technical work and we invite others who are interested in DNS and have expertise in this area to join what we call a caucus and collaborate and work together to delve into deeper matters. And these members are appointed by RSSAC, caucus members. Next slide, please.

Today, in all, we have 88 caucus members and this includes the RSSAC members as well. And each of us has public statements of interest and they get credit for work that they’ve done. And if you go to our website, you will see that.

And the purpose, as I said earlier, is it’s a pool of experts. We lean on them for expertise. They create a critical mass of subject matter experts and their skillset is very broad. It’s a broad spectrum of individuals and there’s transparency of who does

what. It's who, what expertise and what are the hats they wear. And essentially, it's the framework for getting work done. So we're results-driven. We have very clear, deliverables for every work assignment. We assign leaders and deadlines.

And we invite anyone who is interested in being a member of this community to send an e-mail to RSSAC-membership@icann.org and we do have a membership committee that reviews your application and then will get back to you. Thank you. Next slide, please.

The next meetings, caucus meetings of this year, one in July, ATF 102 in Montreal, Canada and we meet at every even-numbered IETF meeting and every – the AGM. Correct? If I'm not mistaken, it's every AGM meeting of the three ICANN meetings so the next upcoming one is in July 2018 and then ICANN63 in Barcelona, Spain in October. Next slide, please.

We've also recently created an FAQ and what you see in front of you is a URL to the FAQ and just frequently asked questions that come up over and over again. So we've put this out there for public consumption. And as I said earlier, we've got five liaisons, so three of them are coming from other organizations into the RSSAC and two go from the RSSAC caucus out into the community and they serve various purposes.

So one of them is Russ Mundy and he is an incoming liaison from the SSAC into the RSSAC. And I'm going to ask Russ to say a few words about his role here.

RUSS MUNDY: Thank you, Tripti. So the SSAC is the Security and Stability Advisory Committee, another advisory committee and it is similar but different in many ways, similar to RSSAC in a number of ways, different in many ways. And one of the biggest, obvious differences is scope where the RSSAC is focused on the root server system, the SSAC scope is essentially Internet-wide with the most important things that get attention that impact that center of the infrastructure. So you could say that the SSAC is a little broader than the RSSAC but they don't go as deep in certain ways.

So what happened where we really established the liaison role, I think formally, was around the time that the root zone was being initially signed. There had always been some informal coordination and back and forth between the two, but we assigned a formal liaison at that time because clearly, there was a very large interest in DNSSEC from the SSAC perspective and the RSSAC and the focus of the RSSAC work had to make sure that DNSSEC would work in the realm. So that was kind of the genesis point, I believe, where the liaison was first assigned.

And we found it beneficial from the SSAC side and the RSSAC, I think, has too. And so we do share information back and forth and keep each other informed about what's going on. Thanks.

TRIPTI SINHA: Thank you, Russ. The next individual I'd like to introduce is Daniel Migault and he's another incoming liaison from the IAB into the RSSAC. Daniel, could you say a few words please?

DANIEL MIGAULT: Yeah. So the IAB is the Internet Architecture Board and its main focus is to look at what could impact the Internet architecture. The IAB is closely working within the IETF community which is mostly looking at the protocol part of that. So as a liaison, I am making the bridge between these two communities which means informing the technical and the architecture community about the work being done at RSSAC and basically, their interest is to see if some of the discussions lead to a major impact on the architecture of the Internet, so we have our discussions and I am doing the shuttle between those two communities. It's relatively light but it enables to have both of us, both communities, aware of what each other is doing and to understand, also, each other.

TRIPTI SINHA: Okay. Thank you, Daniel. And now I'd like to introduce Naela Sarras. She is our incoming liaison from the IANA Functions, which is now called the PTI, and I would love if you could introduce yourself and say a few words.

NAELA SARRAS: Thank you. So I am Naela Sarras. I am the liaison from PTI to RSSAC. PTI happens to be the organization that currently houses – so PTI is Public Technical Identifiers. It's an organization that was founded right after the transition and it's an affiliate of ICANN. It currently houses the IANA functions and that's why we refer to ourselves here as the IANA Functions Operator.

So as the IANA Functions Operator, we have three major areas that we work on, protocol parameters, number resources, and domain names. And our role here falls under the domain names area.

Our interests and our liaison relationship to RSSAC is because of our role in operating the root zone. So we're the entity that is taking requests to either modify or add entries to the root zone and the root server operators are the ones that are serving that root zone that we're helping generate.

So as the case is with any other zone operator, it's important for them to have a close relationship with their name server

operators and that's exactly why we're here, to listen to their deliberations and see how, as they think about where we're headed as an industry, how the Internet is evolving and they're thinking ahead, our role is to think about, well, how does that impact what we're doing today and how does that inform what we're doing today and might change it.

TRIPTI SINHA: Thank you, Naella. I earlier said that we have five and I realize that we are going through an older deck of slides and actually, I would be remiss if I didn't continue to introduce three others who are actually closely tied to what Naella does.

One individual is Duane Wessels who is actually the root zone maintainer and he had to catch a flight so he's not with us in the room, but that, essentially, is his function. And I'd like to now say my Co-Chair, Brad Zerc. Brad, I'm sorry.

BRAD VERD: Thank you.

TRIPTI SINHA: I said Zerc. I'm sorry. Brad Verd is our liaison to the RZERC. Brad, could you describe your role?

BRAD VERD: Yeah, I'm an outgoing liaison to the RZERC Board which is the Root Zone Evolution Review Committee, so I represent RSSAC's interests in that role. That committee, yeah, there's been no action on that committee so we are quiet.

TRIPTI SINHA: Okay, thank you. And Liman is our liaison to the Customer Standing Committee, so Liman.

LARS-JOHAN LIMAN: Yes, Lars-Johan Liman here. I am the liaison to the Customer Standing Committee, or the CSC, which is also a new body created after the transition of the stewardship of the IANA function and it is basically an auditing function of the operations of the IANA function. And that is also a very small group with a very focused task and there, too, I can't tell you.

It's an ongoing effort but it's very smooth work, very much thanks to Naella and the team at IANA. It just works.

TRIPTI SINHA: Thank you. So moving on to, pardon me?

UNIDENTIFIED MALE: [Inaudible]

TRIPTI SINHA: Oh, Fred? I'm not done yet. I know there are two more. I've got more. I was just going to say we have another liaison who is outgoing from the caucus to the NomCom and I'd like to introduce Alejandro Acosta.

ALEJANDRO ACOSTA: Well, hello. I am Alejandro Acosta and I am the liaison for NomCom, and basically, more or less as Daniel said, what we are doing, we are the bridge, I am the bridge between RSSAC and the NomCom. Basically, what I try to do there is to represent the technical part with SSAC and maybe the IAB. We are a non-voting member and I wanted to say that this is my last term and probably you need to look for another candidate to put there.

In the NomCom, what we do is to try and fill leadership positions, is what the way it is called, inside the ICANN organization.

TRIPTI SINHA: Thank you. And last but not least is the individual on my left, Kaveh Ranjbar. He is our liaison to the board.

KAVEH RANJBAR: Thank you very much, Tripti. So yes, for the, by definition, [inaudible] liaison from RSSAC to the ICANN Board is basically a courier, sending messages between, mostly from RSSAC to the board. And that's a normal task. And actually, it happens quite frequently, like I think every other month when there is a new appointment and RSSAC needs to inform the board.

But then there are some other which are not documented but are expected. For example, in general, there is an expectation in the board that the RSSAC liaison can answer questions about DNS [readily]. Most of them are basic technical questions so they are not controversial. It's just how DNS works in that realm of technicalities of DNS, normally not things that needs debate.

If matters are of such proportion that there are difference between technical people on how things work, then that's something that normally I bring back to RSSAC and get RSSAC's opinion and share with the board.

And another part is, basically, coordinating the expectations. That's mostly actually liaising from the board to the RSSAC. So for example, when we know something is going to happen, there is an expectation to form an opinion, I try to bring that back as soon as possible and to prepare RSSAC to form an opinion and advise.

Finally, because as an outgoing liaison, you would become a board member, the board liaison is not a director but still a board member. There is expectation to do regular board member duties, which includes being part of board working groups and board committees. For example, I am part of Board IDN Working Group, Board Risk Working Group and I chair Board Technical Committee, Board Risk Committee and chair Board Technical Committee. So that's basically the gist of it.

TRIPTI SINHA: Thank you, Kaveh. So this liaison framework that we have is to facilitate communication with these different communities as well as to ensure that where our expertise is required is given and where we need expertise, it's brought in. So with that said, I think I'm done with my part. Next slide.

I'm going to turn it over to Brad now. Brad Zerc.

BRAD VERD: Thanks, Tripti. Brad Zerc. Thank you. I will just add to the liaison portion that we are open to other liaisons from any of our constituencies should they want to engage with us. We don't limit them. We're open to them and if there is a liaison that is not represented here, it's because that group has either chosen not

to engage or there isn't really necessarily a need. So I just want to add that.

Moving on to the next, we're going to talk about recent RSSAC publications since the last ICANN meeting. We have two. Next slide, please.

We have RSSAC 30 and RSSAC 31. RSSAC 30 is, we talk about the root sources and where the entries are that define a root server operator, and then we have a response to the GNSO PDP. So next slide.

All right, RSSAC 30, this was a short, simple statement that, for us, we believed was kind of like a foundational statement, really to be used. We realized that this was not called out anywhere within ICANN in the policy world that says, "Where is a DNS root server operator defined?" So RSSAC made a statement saying that root server operator is defined in these three files, the root hints file, the root zone, and the root-servers.net zone. And we made this in the expectation that we will be using this as a reference going forward with some of our work.

Any questions or comments about that? It was a rather noncontroversial statement, kind of just stating history. No? All right. Next slide, please.

RSSAC 31, this was a response to the GNSO policy development group that reached out to us asking about the limitation on delegations on a per annum or per annual basis. So you got to roll back in history here. You got to go back to the pre-new GTLD days. There was a root scaling study done back then that, essentially, gave a recommendation that limited the number of new TLDs added to the root to a thousand per year. And that, obviously, drove any number of different things on the logistics side for IANA, for ICANN, for a number of different groups. As the new TLDs came out, there was this kind of looming number of a thousand that the recommendation was to not go over.

So the question from the GNSO was, "Is there a new number? Can we change that number?" and there were a whole bunch of numbers thrown out from different groups, anywhere from, well, I heard up to about 20-some-odd thousand per annum. And really, what happened after, when RSSAC sat down – can we go to the next slide? When RSSAC sat down, we had a long discussion about this and we came to the conclusion kind of quickly that it's really asking the wrong question. Rather than asking how many can we add per year, we thought it was more important that it should be the rate of change per year and then we put kind of an absolute magnitude on it saying that it should be kept under 5% per month.

We considered the root zone to be a – how do I say this? – it is a TLD but it is not any TLD. It is the root zone, so changes to the root, modifications, should be made in a diligent fashion, very carefully, very methodical and the idea of adding 25,000 names to us at any given rate was just not a safe practice.

So we, again, our recommendation was to rather than focus on an absolute number, focus on the rate of change and we limited the absolute amount of change to a 5% per month and we shared a number of different things as to why we came up with that. Obviously, stability and resiliency were our main goals. And we talked about history as conservatism towards these types of record resources allowed in the root zone is appropriate, but again, it's about safe and methodical change. Next slide.

We also briefly stated that if and when a New gTLD Program goes into account, that as change is being introduced to the root zone, that there needed to be a clear and methodical plan for backing out should something be identified as creating a problem. How do we back that out? How do we implement it back to a good known state?

And then also, we echoed both SSAC's recommendations and previous recommendations from the root scaling study back in 2010 or 2009, something like that, where we said an early warning system for the root server system should be developed.

That has not been there yet and that recommendation still stands.

There were some other comments that kind of came out. They were questions. These were not part of our official response, but there were questions that were asked of the GNSO when we sent it back to them and one of those was the idea of creating n number of TLDs. Does it flatten the name space and push traffic up to the root servers? Something to think about.

And then, again, going back towards the how to proceed, we referred back to the CDAR study, which was the reference of the GNSO and their question that the CDAR study said that the rate of change to the root zone should remain gradual. So that was our response to the GNSO and those were the two publications that we have completed since ICANN60 in Abu Dhabi. Any questions?

[ABDUL KARIM]: Thank you very much. My name is [Abdul Karim]. I'm from Nigeria. I'm just wondering why you are worried about traffic going back to the root servers because I think the reason why we have so many DNS servers now around the world is because to [get there for] more traffic and this is like growing every time. So why is that a worry about having traffic going back to the root servers? Thank you.

BRAD VERD: I think you're referring to this creating a flattened name space comment. This was raised by RSSAC as if, yes, DNS traffic is growing, continues to grow, but at the root level, as we've stated, most traffic is seen at the resolver level and cashed. By flattening out the root zone or the actual name space, I should say, the question – and this was a question, this wasn't a statement – the question that maybe we need to look at is, "Does this, by flattening out the name space, pull traffic closer to the roots and would traffic levels increase at the root server system?" And if so, obviously, there are, that needs to be conveyed to the root operators so they can account for growth and change.

Anybody else want to add any commentary to that?

ABDALMONEM GALILA: Abdalmonem, ICANN Coach. Do you think that using local root zone at [ISPs] will solve the issue of the traffic at [real] root servers?

WES HARDAKER: That's a good question. So it actually turns out for things like… Yes. So his question was if you used a technology like local root, and he specifically mentioned my project which he was at my

presentation the other day, but it's also been talked about by the board lately as hyper-local root service. That's really keeping a copy of the root zone in your local resolver.

That works really well and it would feel like, since you're not sending any requests to the root, that you would be generating less traffic. The reality is that you're actually probably going to use more traffic because you're pulling down more data than you're likely to use. Nobody goes to all of the TLDs. So you may go to, say, ten TLDs during a day when the reality is if you pull it down to your local copy, you're going to pull all the TLDs. You're actually pulling more information than you're likely to use. So no, that actually, unfortunately, doesn't solve the bandwidth issue. Good question though.

UNIDENTIFIED MALE: Are any TLDs likely to be removed? [Inaudible]

Yes. What is the story with TLDs disappearing? I think I've seen one go away over the last two years. Are any going away for abuse or what reasons are there for TLDs going away?

BRAD VERD: That's a great question. Unfortunately, it is out of scope for RSSAC as we run the root server system that serves the root zone, and obviously, we serve the zone as given to us by the

IANA function. So policy questions around what gets removed and added goes to a different group.

I know the ccNSO is currently going through a PDP to define how to remove a ccTLD and there are lots of processes in place to remove TLDs and gTLDs, all of which, I believe, have happened, removal of them.

Yes, Michael.

MICHAEL CASADEVALL: Out of curiosity, what is the expected impact in terms of bandwidth of the changeover from, of the KSK changeover since it's going to increase a larger key size, and thus, increase the size of the RrSG records from the root?

BRAD VERD: Thank you for the question. Again, that question is a bit out of scope for RSSAC and this meeting. It is a question that probably for discussion within the caucus and/or… Go ahead, Russ.

RUSS MUNDY: Tomorrow, Wednesday, there's a DNS workshop that will be a whole room full of people that have looked at lots of different aspects of DNSSEC including the KSK rollover. But right now, both of the keys, the current or old key and the new key, are

actually in the zone and being served right now. So I don't know that there's been explicit traffic size and data studies on it, but there has been no known reports that I'm aware of, anyway, of negative impacts of having both key sets in it.

BRAD VERD: I will add that, as Russ, stated that all the keys are out there. They've been out there for a long time when this got put on, when the actual roll of the KSK was put on hold. So we've been running with a much larger response size for an extended period of time now and there was a lot of study that went in prior to the publication of those keys because we knew exactly what the size of the responses were going to be at – I think there were like three different dates. There was the introducing the new key. There was signing the new… There was a whole bunch of different steps and, obviously, based upon the result and the data collected out of RFC 8145, the decision was made to hold off on rolling the key and OCTO has a plan to move that forward.

All right, no other questions regarding the recent publications. We will move on and that brings us down to updates on current work, which I am going to turn over to the gentleman to my right, Mr. Liman.

LARS-JOHAN LIMAN: Lars Liman here. So we have, I'm going to report on two ongoing work parties. The first one where I'm the shepherd, the liaison, the caucus regarding anonymizing DNS query sources. The root server operators collect a lot of statistics and some of it is put in central repositories for research and other activities.

And as we do report the entire incoming query, it has both a domain name that's being queried for, but also the source IP address from where it was sent and in some circumstances, that can be seen as an integrity-sensitive combination of information. So some of the root server operators, when they upload information, they anonymize the IP address of the sender so it cannot be traced back to the original computer. But that makes it a bit difficult for the researchers to do research in a good way when they combine information from various operators because the same IP address will not appear in the same fashion.

So this work party is tasked to look at, for instance, to see if there is an optimal way to do the anonymization and whether it is something to recommend that all the root server operators do. And if they arrive at the conclusion that this is a good idea, also to recommend a specific way of doing it so that all the root server operators do it in the same way.

This work has come fairly far along its way. There is a draft document. It has still a number of comments that need hashing out, discussion in the work party. And I believe also that they are trying to find or put together some test code to run to see that the ideas they have work properly before submitting this as a final proposal. So that's that one. Next slide, please.

The next one is on packet sizes. As was discussed here with the DNSSEC keys and stuff, the modern DNS packets can turn to become fairly large, especially those that contain keys of size and if you have multiple keys and so on. And that makes the server responses going back start to touch on various settings in the systems along the way. It's both in the DNS software, it's on the actual computer platform that we ran, but also in the nearby rooters.

We have tasked this work party to look at the various networking parameters to see if there are any recommendations in order to make the traffic flow as smooth as possible. So there are a number of deeply technical details in networking that they are being asked to look at. This work has not really just started, but it's in its initial phase so there is still a lot of work ahead in this work party.

Any questions on those? Otherwise, yeah.

ABDALMONEM GALILA: Abdalmonem, for the record, ICANN Coach. I think if root server allows zone transfer, is it secure to do that? If yes, why wouldn't do that for the other root server?

UNIDENTIFIED MALE: We do allow free zone transfers. I believe they're secure. I don't know of any reason they wouldn't be.

LARS-JOHAN LIMAN: I don't either.

UNIDENTIFIED MALE: So I believe they're secure.

ABDALMONEM GALILA: How come secure? How come secure? If I am an [inaudible] tech and I want to see the keys or see the details of that root zone file, I will take it as my guidelines.

UNIDENTIFIED MALE: So first of all, multiple operators allow zone transfer. It's not only [F]. I think, I don't know the number, but I think 8 or 9 out of 13 actually do that. That's first.

And second, root zone is a public file. It's openly published. You can go to IANA's website. You can even download it over web, so

you don't even need to download it over zone transfers. Unlike many TLD zone files or ccTLDs, some of them actually treat that as business secret, but root is not like that. Root is completely open, has always been open. There are even databases which keep track of root zones so you can't even see all the versions of root zone. So it's an open file and it's signed so it's not even in a cryptic. It's just signed, so normal root zone. It's a normal zone file.

BRAD VERD: I will add that it's also available and has been available for I think the beginning of time on RS.Internet.net, which is currently run by the root zone maintainer.

WES HARDAKER: Yeah, I have the list of letters in front of me that support zone transfers and includes B, C, F, G, K and then ICANN actually has two servers that they support it from. Note that the transfer of it is if you want to do a secure transfer to guarantee that you're getting the source, you should do it over https or with a TSIG-encrypted or TSIG-protected DNS request, like local group provides, like you were talking about earlier. If not, you had better verify it with DNSSEC if you're not going to use a secure form of transport.

LARS-JOHAN LIMAN: Thanks. I'll hand it back to, I guess, the next.

UNIDENTIFIED MALE: Regarding the study of the anonymization, do you have a deadline of when that's going to be like a final work release?

UNIDENTIFIED MALE: No, we don't have a deadline. No. I believe that it will be finished before the end of May at least. It's near conclusion.

BRAD VERD: I think the next slide, please, is a tool slide. Wes, were you going to speak to this?

WES HARDAKER: Yeah, just a little bit of history. We decided that we should create a single place to house tools that were related to RSSAC Caucus work efforts and, specifically, or even just the root server system generically.

And so, we have four tools today that are available on that GitHub page which is GitHub.com/rssac-caucus and there's all of the RSSAC002 data is housed in a repository so you can just pull

that rather than having to go collect it individually from each root so it's a little bit more convenient.

There's an R-API for parsing that RSSAC002 data. There are the root server naming tests that were used in producing RSSAC0028 and speaking of the anonymization harmonization effort, there is a new repository there that's actually new since the last time we met and that has some pretty heat map pictures if you want to go look at them.

Oh, a quick follow-on. So if anybody does have tools or things that you're developing in order to deal with the root server system in any way, feel free to ask us for a repository. It's not limited to just official work efforts under RSSAC. In fact, the R-API, for example, was a code that was written outside of an actual work party effort or anything. We're trying to build a catalog, an easy place to go find things, especially for RSSAC002 data, but just for the root system in general.

TRIPTI SINHA: Do you want to continue with the rest of the slides, [inaudible]?

WES HARDAKER: Is it on to my section next?

TRIPTI SINHA: Yeah. Next slide, please.

WES HARDAKER: Oh, you're right. It is. Okay.

All right, so the next, the final section, I think, today is about community interaction so go ahead and go to the next slide.

So first off, the RSSAC is the policy and sort of high level thoughts and management of the root server system when it comes to the root server system, this RSSAC and the RSSAC Caucus within ICANN are where sort of high level decisions are made about the root server system or evaluation.

We're working very hard to ensure that… transparency is one of our goals and we do, unfortunately, have to deal with sensitive information some of the time so we do have some closed meetings. Sensitive information might include things like infrastructure and service details or attack and vulnerability analysis that can't be released immediately. So some sessions, unfortunately, have to be closed. We do try to be transparent whenever possible and we try and think about every topic and what, as new topics come in the door, whether it contains sensitive information or if it maybe still contains sensitive information so we can go from closed to open in some cases.

And generally, we try and house all work that is possible under the RSSAC Caucus. So again, if you're interested in helping out with the RSSAC, that is where we try and produce most of the documents that we can. Let's see, so this list in particular is a list of some sources of information about RSSAC. So all of our meeting minutes, for example, are published under that top link. All of our publications are under the second link. They're all available for download.

Our Root Server System tutorial, which we give twice at every ICANN meeting is available. The latest slides are there. And then if you're interested in diving into how RSSAC, as a body, functions, that last one, RSSAC00 Version 3 is the current version of our operational procedures and sort of acts as our bylaws. And next slide, please.

On the flipside, the folks that deal with the technical day-to-day operation of running the root server infrastructure are the root server operators. They deal with things like devices, and cables, and connections and configurations and stuff like that. Some people that work in RSSAC happen to also do the technical side, but many root server operators have a completely different set of people with technical skills versus those with skills for ICANN. So you'll find that the two groups do not overlap completely.

The group of technical operators is also making efforts to be as transparent as possible, although, they, of course, have strong needs to discuss sensitive technical issues as well. This is a list of the information that the RSOs publish on a regular basis, so they have a, they publish the agendas of the root server operators and that's on the root-servers.org site which houses most of the information. There's a public webpage that contains news when announcements are made about events that are seen or things like that. Those are also posted there.

There's collaborative reports on major events that I mentioned a second ago. The RSSAC002 statistics are another way that the root server operators are trying to be transparent and release consistent statistics information about each of the root server instances when possible. I think everybody is now delivering RSSAC002 statistics, I'm happy to say.

And then each of the root server operators typically has individual web pages. I don't think we're missing any now. I think, does everybody have one? So all root server operators do have an individual page as well.

Somebody is not on mute.

There's some public letters with IANA and then RSSAC can correspond to technical questions about the root server system. There's an Ask-RRSAC@icann.org e-mail address, Ask-RSSAC,

that is. And the purpose of that mail message was sometimes the community might have questions about the entire root server system that are technical in nature, that you really need to go talk to all of the operators at once in order to find common bits of information, and RSSAC has agreed to be sort of the front end of that so that we can sort of compile that information and get it back to you after it's been collected. Next slide, please.

Most importantly, we're always interested in more information about transparency and so we're looking feedback of how are we doing. So the RSSAC liaisons is another effort, by the way, in order to promote transparency so, as I think Brad said earlier, we're always open to other constituency liaisons where there is sort of a need for cross-communication or engagement. We don't have them with every member because sometimes we've determined that, both sides have determined that engagement isn't actually necessary.

But we're always looking for more and we're always looking for more feedback about our efforts so these questions in particular are really designed to highlight that. Are you aware of all these things that we're publishing and how we're trying to be transparent? And is there anything missing from these lists that you think we should be transparent about that we haven't considered? And what else can we do to improve our own transparency? Feedback is always wonderful.

And with that, I think the next slide is actually about questions. So there's some informational sources there if you want more information about RSSAC. Again, the main webpage is RSSAC.ICANN.org and our publications are fascinating reading, especially at night right before bedtime.

There is an FAQ which is well worth reading. We get a lot of questions that are common, especially during the tutorials and we've tried to outline sort of the most common questions with answers in a consolidated place. And then finally, the RSSAC Caucus has its own webpage for the work that it's undertaking.

Anybody have questions about anything? And I'm going to turn it over to Tripti and Brad at this point.

BRAD VERD: Any questions?

TRIPTI SINHA: Hearing none, I guess that this brings us to the end of the meeting. Thank you very much for coming and the meeting is adjourned.

BRAD VERD: Thank you all.

[END OF TRANSCRIPTION]