ICANN78 | AGM – RSSAC Caucus Meeting
Tuesday, October 24, 2023 – 10:30 to 12:00 HAM

OZAN SAHIN:

Hello, and welcome to the RSSAC Caucus Meeting. My name is Ozan, and I will be the remote participation manager for the session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior.

During this session, questions or comments submitted in chat will only be read aloud if put in the proper form, as noted in the chat. I will read your questions or comments aloud during the time set by the chair or moderator of the session.

If you would like to ask your question or make your comments verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor. Please state your name for the record, and speak clearly at a reasonable pace. Mute your microphone when you are done speaking.

The session includes automated real-time transcription. Please note, this transcript is not official or authoritative. To view real-time transcription, click on the Closed Caption button in the Zoom toolbar.

To ensure transparency of participation in ICANN's multistakeholder model, we ask that you sign in to Zoom sessions using your full name. For example, a first name and last name or surname. You may be removed from the session if you do not sign in using your full name.

And a last reminder that I will be circulating an attendance sheet. So if you can, just circle your name on the attendance sheet, and pass it on to the other participants. That will be appreciated.

With that, I will hand the floor over to RSSAC chair, Jeff Osborn.

JEFF OSBORN: My Zoom mic is open. That's a mistake. Good morning. More coffee, please? Sure.

Good morning, everyone. My name's Jeff Osborne. I'm with ISC. And as part of calling this meeting of the RSSAC Caucus together, I think we should do introductions around the table. So let's start with Wes.

WES HARDAKER: Good morning, everyone. This is Wes Hardaker from the University of Southern California.

ERUM WELLING: Erum Welling from DISA G-root.

BASTIAAN GOSLINGS: Bastiaan Goslings from the Public Policy Team of the RIPE NCC.

BRETT CARR: Brett Carr, AWS.

KIM DAVIES: Kim Davies, IANA.

ANDREW MCCONOCHIE: Andrew McConachie, ICANN staff.

DANIEL MIGAULT: Daniel Migault, IETF liaison.

JEFF OSBORN: Jeff Osborn, ISC and RSSAC chair.

KEN RENARD: Ken Renard, Army Research Lab, RSSAC vice-chair.

KARL REUSS: Karl Reuss, University of Maryland, Root Server Operator.

DUANE WESSELS: Duane Wessels from Verisign as the Root Zone Maintainer liaison.

HIRO HOTTA: Hiro Hotta, [inaudible] and JPRS RSOs.

ROB CAROLINA: Rob Carolina, ISC.

BRAD VERD: Brad Verd, Verisign.

JEFF OSBORN: Thank you all very much. It seems to be awkward trying to do that with the remote people, so I think we'll leave it at that. We have an hour and a half scheduled today, which strikes me as long for some of the items we've got. We have a little housekeeping on the part of the Caucus, and then reports on ongoing work, and then the Ask an RSO section at the end that I think should be interesting with Ken hosting.

So to begin, the RSSAC new Caucus members. Ozan.

OZAN SAHIN: Oh, yes. Thank you, Jeff. So I'm just checking if we have the new RSSAC Caucus members in the meeting room. And what I mean by new Caucus members is the Caucus members who joined since June 2023. I believe we have Joe Hayes in the room. So, Joe, if you don't mind introducing yourself, coming to the mic, that would be appreciated.

JOSEPH HAYES: Good day to everyone. Joseph Hayes, Caucus member from DISA.

OZAN SAHIN: And I believe Joseph is the only new Caucus member that we have on the call.

JEFF OSBORN: Okay. Like I said, this could be a pretty simple meeting. The next item, I believe, was still part of the Caucus Membership Survey which we haven't formally presented here. Again, Ozan on stage.

OZAN SAHIN:

Yes. Thank you, Jeff. Hello, everyone. This is Ozan from support staff. So as you may recall, the annual RSSAC Caucus Membership Survey was conducted in September, and we would like to discuss some of the results here in this meeting.

Before going into that, let me share the RSSAC Caucus membership data as of today. We have 120 RSSAC Caucus members, and an important part of the members are from North America. This is followed by Asia region, almost by 30%—to be exact, 29%. And then we have, also, members from Europe, Africa, and Latin America.

So going back to the annual survey, this year we received 36 responses, and this is more or less in line with how many responses we get each year. There was a decline in 2021, but this is not surprising.

And one question we asked is whether the respondents participated in an RSSAC Caucus meeting before. So there's a decline this year if you look at the "yes" response. I believe some higher portion of the respondents were the Caucus members who haven't been to an RSSAC Caucus meeting in comparison to previous years. But still, 64% of the respondents participated in an RSSAC Caucus meeting before.

One question we asked was what the respondents thought about meeting frequency and whether RSSAC Caucus needed to meet more frequently, less frequently, or the current setup is fine. And the most popular answer that we got for this question on each of the previous three to four years was keeping the current sequence.

But this year, there was a slight decline in this. In the responses, 69% responded supporting keeping the current sequence. And if you look at

EN

the responses supporting having more frequent meetings, or much more frequent meetings, especially in the more frequent meetings area, you see a slight increase there.

And the other question was about the meeting venue. Currently, RSSAC Caucus meets at ICANN Annual General Meetings, as we are meeting today, and also even-numbered IETF meetings. That is what is referred to as keeping current venues. Again, this is the most popular answer with 81%.

Another question we asked was whether the respondents participated in a work party, and this is more or less in line with the responses that we received in the previous years. 81% of the respondents participated in a work party, whereas 19% did not.

Another question we asked was whether the respondents contributed to a work party and how they contributed. So most of the respondents confirmed that they participated in discussions. 86% did that. And also, 66% of the respondents confirmed they reviewed text previously, 31% indicated that they wrote text, and 14% of the respondents said they served as work party leaders.

And we also asked how difficult it is to contribute to an RSSAC Caucus Work Party. And the responses in this question were a bit mixed. The most popular response was 40% finding that contributing to RSSAC is easy. And then 33% were neutral. 24% found it complicated to contribute. And only 3%—or I think that corresponds to one person or one respondent—found it very easy.

78 | ANNUAL GENERAL MEETING

And we also asked what could be the reason of non-participation in RSSAC Caucus Work Parties and suggested some potential reasons. So in this department, the most popular response is "there's no time to contribute" with 39%. And 28% said they had no technical experience. Again, 28% indicated other reasons. And 22% said they had no interest in the topic. And 6%, or two respondents, said there were communication barriers.

So one small change in the question of whether the respondents wanted to continue as an RSSAC Caucus member. In the previous years, all of the respondents wanted to continue. But this year, we had one respondent who wanted to resign from RSSAC Caucus, whom we removed after the survey.

Also, we received some free-form text responses and comments. So, of course, the results were reviewed by the RSSAC Caucus Membership Committee and the RSSAC Admin Committee. And these committees wanted to highlight some of these free-form text responses.

So one interesting response in response to the question, "Why have you not participated in a work party" was that the current Caucus members were there for many years, and they knew each other very well. And this was demonstrated as an advanced and cohesive group, and the bottom line is that the group didn't need any further contribution. So that was the feeling the respondent had.

And also, the other feedback to highlight. So some form of mentoring was suggested for the participation of new people. And I think upon discussion within the RSSAC Admin Committee and the Caucus Membership Committee, a potential suggestion to address this

comment was kind of onboarding Caucus Work Party meetings to provide mentorship to new Caucus members, or new Caucus members that are attending a work party for the first time. So that was a suggestion.

Another feedback that we received was making the sessions interesting with multimedia and graphics. Currently, according to this respondent, they were very academic. And one final feedback that we'd like to highlight suggested that partial funding could be given to more RSSAC members for attending meetings like IETF ICANN in person.

And finally, there were some suggested items for RSSAC Caucus to study in response to another question in the survey. We also highlighted these suggestions. So on the agenda today, there's a discussion for upcoming work, if I'm remembering correctly. And we also wanted to discuss these suggested items in this meeting. We can do it now, as we're talking about the findings.

But if you want to put the focus now on this or other findings, or whether we should go ahead and start the upcoming potential RSSAC Caucus work in light of these suggestions, I'll defer to you, Jeff.

JEFF OSBORN: I had my hand up on something else, so I think this part should be later. It makes sense. But if you can go back a slide.

There's an ongoing conundrum. I'm the chair of the Root Server Caucus Membership Committee as well as the RSSAC, and I'm attempting to step down for the next year, if somebody else would like to come in and do it.

One of the things that is challenging is there is a conflict in mindset, although the written direction is pretty obvious. The purpose of the RSSAC Caucus is to gather a group of experts who can be helpful in getting work done for the members of RSSAC who are lacking in areas of expertise.

Now having said that, that's a big call because there's a lot of depth to the members of RSSAC in terms of technical expertise in operating DNS in very large scales. So what we often get is somebody saying, "I'm a student. I am new in this business. I would love to burnish my credentials and learn." Where in other parts of ICANN, the idea is to cast a wide net and be very inclusive, we've, from a rereading of the rules, come to realize the Caucus is not that place.

So we've tried to be a little kinder about it where if you apply and you do not have what we consider to be a wealth of hands-on DNS experience that could be useful or other skills that we might be lacking, you get a polite letter saying you might want to show us further what it is you think you bring to the party.

And we also offer to put them on a mailing list where the information has always been available on a pull basis, but we put it up as a push, and you show up on a mailing list.

This has been part of a process we've been working through. It was interrupted by COVID where people who don't show up or attend meetings for some fairly long period of time are gently told, "Thank you for your service." You know, "Your services are no longer needed. Thank you for them."

So I'm just trying to put that in place because very often at ICANN, we're all about how can we be more inclusive, and how can we add more people? And the RSSAC Caucus is very specifically trying to be an expert body, so it's very difficult to be broadly inclusive and expert. And that's something we struggle with most months.

If there are any questions, this is where we'd have them. Wes.

[WES HARDAKER:]   [Erum was first].

JEFF OSBORN: I'm sorry. Erum.

ERUM WELLING: Thank you. So you raise an interesting point about the purpose. I'm sure it's in a Charter somewhere. So it seems like there is a need, though, for this inclusive group to exist so people feel ...

Perhaps there is something brilliant in their minds that needs to come out on the table. Maybe this is not the forum. But if ICANN is the venue for a multistakeholder community to come together, then perhaps there needs to be another group. I think it's great that we're telling people if you don't have anything to bring to the table, you may want to reconsider, but I think there's also another side to that of being inclusive in the sense of having suggestions.

I would prefer that at the end of our discussion that we just had, that somebody would say, "That's fine. Maybe this is not the right

environment." But it does seem like there is a desire to have such an environment, and maybe we need to brainstorm about what could be done to make these people feel more inclusive. So thank you.

JEFF OSBORN: Thanks, Erum. The reason I bring this up almost every single opportunity I get is because I feel that conflict myself. I run a technology company. And when you have yet another brilliant, experienced 54-year-old white guy who used to work at the same three places the last three hires did, you have to remind yourself that it starts here that you gain a broader group of people to discuss.

We kind of ran into the Bylaws and saw that they were really looking for expertise and sort of had the come-to-Jesus about a year ago of, "This is not something if somebody's trying to burnish a CV or add their education." And so we bring it up. I'm not sure who's going to say, "Hey, go make something more inclusive," but it's certainly not for lack of trying. So thanks for adding your voice.

Daniel.

DANIEL MIGAULT: So I think there is probably middle ground between all that, which is, I suggest that if someone is new, we may ask him to review the document. And then the chair may take a little bit more attention, when it's a new person, to the comment they're being provided on the document. I think that's probably a reasonable alternative versus seeing, "Okay, so here is the DNS 101 lecture I'm going to give to the newcomers."

So people can get involved. We're helping to get them involved, but we don't commit too much so that we don't repeat [inaudible]. But on the one hand, we become more inclusive. We address the purpose of if we're a known group, someone can join, but we are not slowing down the work party. So it's just a matter ...

I think we don't need to re-architect the way we work. I see the demand as very simple, as maybe the chair is taking a little bit more attention to the comment being provided in a document. That could be one way to do. And I don't think the overhead is too high for that.

JEFF OSBORN: Like I said, I think that the problem exists. I think the solutions are not yet obvious, and everything is complicated in a largely volunteer organization. So that's the problem set. If somebody wants to come along and fix it, I guarantee you'd be unanimously elected to fix it.

Wes.

WES HARDAKER: This is sort of slightly tangential, going back a bit ago in the conversation. But one thing that just occurred to me is we have a lot of Caucus members that may not be participating in the last year. And the reason being that their DNS experts are their experts, and this is their space, but the work parties we had in the last year were just not relevant to their particular sub-skill set within the expertise skill set that we recruit.

So that might be something we want to put in future surveys or something like that. "Why didn't you contribute?" And maybe offer them that suggestion that "my expertise didn't align with the work parties this time."

Like, right now, the only one that we have open is a Security Incident Reporting Work Party. Well, if you've never done operations, you've never done security reporting, and you're an expert in the protocol bits, then you're probably not going to help much in that work party, or may not want to help much.

JEFF OSBORN: Good point. Anyone else? Online, Abdulkarim.

ABDULKARIM OLOYEDE: Thank you very much. Apologies for not being in the room. I just wanted to say that the working on DNS, I totally agree that, yes, there has to be some level of expertise or technical expertise that has to be brought to the table.

However, if one of the requirements would strictly be working on a DNS server, then I think a lot of people will have never and would never have the opportunity of working on a DNS server because we all know it's not something you can get to work on every day. Or it's not something I get to work on everywhere. So if that is also a requirement, then we'll put a lot of people—it will disadvantage a lot of people.

So, yes, I agree that there should be some level of technical expertise. But we just need to be kind of careful so that we do not disadvantage people who are already disadvantaged. Thank you.

JEFF OSBORN: Thanks, Abdulkarim. I see Hafiz has his hand up.

HAFIZ FAROOZ: Thank you very much. Thank you very much for giving me the opportunity. I have a similar observation like Abdulkarim because I'm relatively new, more than a year now so far. And my observation is that although I'm experienced, I've been working in network and DNS since long, but the root server is kind of an area which, even if you're experienced, this is something new for you.

I saw in the recent past there was an ICANN Learn course for the RSSAC, but that [it's also basic]. So maybe, I don't know whether we can have something more advanced in the ICANN Learn or if we can know more about the RSOs. That will give us the confidence that we know your language, [then when] you speak to each other, and then we will be in a better position to comment on what has been discussed.

Yeah, that's my thought. Thank you.

JEFF OSBORN: Thank you. Naveed.

NAVEED BIN RAIS:	Hello, everyone. So I was coming while listening to the meeting because I was about 10-15 minutes late. But one thing that struck me while Ozan was mentioning about the outcome of the survey is a good number of people said that they did not have technical background [or kind of that] to participate.

So I wonder if they did not have the technical background at all related to RSSAC or related to the work parties that were ongoing. So this has to be clarified. Because if they did not have the technical background to contribute to Caucus in any way, then why they are here in the first place?

So I was just wondering if this can be clarified. So I'm not sure if that was—because I was just listening to the audio. But if this was the question, then it was a good number, like 20-something percent said this.

JEFF OSBORN:	That's a really good question. I can only answer for sure for ... How long has it been, Ozan—three years—I've been on the Membership Committee? Two?

OZAN SAHIN:	Your enrollment? Two years, I guess.

JEFF OSBORN:	Two. For two years, I know that when we've had somebody come by and apply, and the normal order is probably six or eight a year, there has been a very high barrier to technical expertise. There has absolutely

been nobody who thought technology was an interesting idea; they should think about it.

Rather, when we say "no technical expertise," we meant their expertise was in another area of either operations or computer design or computer deployment or programming. We have very few people who you would describe as not having any technical experience. It's just a matter of we've been trying to have it be that specific. So I could see where you could get that idea.

How I had been reading "no technical experience" was, for instance, we're doing the security monitoring, which is something I've had no work in ever, and so I've been skipping the entire thing. So I don't think I would bring anything to it. Where some of the other sessions are areas where I've been involved with from the side, and so they're a lot more interesting. There's a lot of self-selecting.

I don't know where or how long the Caucus organization has been founded, but the Bylaws that existed from RSSAC, or from ICANN from way back, do specifically say this should be an expert organization. But the reason I raised my hand was it kind of looked like if you came in and you had good intentions and you meant well and you hope to learn, they'd let you in. And so pointing out the conflict was something we kind of specifically tried to do.

The other thing Ozan was doing that I thought was really helpful was noting where people hadn't attended in, I think it was a year or two, and just saying, "Do you really need to be doing this?" There were a larger number than I expected.

Was it as much as 100 people at one point?

OZAN SAHIN: The number of Caucus?

JEFF OSBORN: Yeah.

OZAN SAHIN: Yes, 120.

JEFF OSBORN: Yeah. So it's a lot of people, and you certainly don't see them in every session. And you don't see them in every meeting, and you don't hear from them in every work group. So the idea of not so much being less inclusive but being more selective, let's say, I think has been a steady one.

So we would love to figure out people who can contribute without a lot of training if they exist in this whole environment. But it's harder to have it be yet another class to try to teach people. Does that make sense? It's a little complicated, and it's a little conflicting.

NAVEED BIN RAIS: Yes. Thanks for the answer. So, for example, for the next round of survey, if we can specifically ask them did they not have the technical background related to the work parties or the discussions that were happening in that current year, that would have a more specific answer, and we would be in a better position to assess and evaluate.

So right now from that text, I'm just guessing. Right? So were they not technical, or were they not interested, or did they not relate to the discussions being held?

JEFF OSBORN: You know what? We try not to change the questions much so that over time, they're consistent, but that's a really good one.

Is there another hand?

WES HARDAKER: No interest in the topic.

JEFF OSBORN: Right. The other thing that would be interesting is looking at the free-form answers that people have at the end of this, and we haven't completely put that all up there.

NAVEED BIN RAIS: This struck me [a lot] because if you go back to the graph, so it says "no technical experience." So it does not say anything else as it is. So it's a big thing if you don't have technical experience. It does not even say "related to RSSAC Caucus" or "related to root servers" or "related to ICANN" even. So that's a big thing. If nobody has a technical experience, then Caucus is not the place. That's what my point was, actually.

JEFF OSBORN: It's a great point. I've looked at these slides a lot of times, and that never occurred to me. It's a really good point, and we'll definitely want to talk

EN

about it because if that's actually what it is and we missed it, we really missed it.

Ken, is this your remote hand?

KEN RENARD: That's a new [inaudible].

JEFF OSBORN: New?

KEN RENARD: New.

JEFF OSBORN: Okay. Ken.

KEN RENARD: All right. So, yeah, hopefully what you're saying is teased out between "no interest in the topic" and "no technical experience." But the people that have been declined from joining the Caucus really have just been "I'm in it to learn." That's the main thing. We're not here to teach people, necessarily, but if you've got something useful, especially as the RSSAC and the Caucus gets into regulatory and political things, those are interesting areas of expertise, as well.

What I really wanted to comment on was Ali Hussain put a good comment in the chat about an onboarding for newcomers to the Caucus. So we do have a Caucus onboarding session that we run, what,

78 ANNUAL GENERAL MEETING

once or twice a year? And that's available to newcomers to the Caucus. But if you haven't participated in a while, or if you just want a refresher, anybody from the Caucus can sit in on those sessions. And we go over things like how a work party works, how it gets initiated, how the publications work, and things like that.

Another thing that struck me in the survey was, you know, people that are experts that have just—maybe they don't have the hands-on keyboard root server-specific ... That's great. We still want your input. You still have valuable things to contribute. So how do you get on board into a work party that is based on a discussion that maybe root server operators have been having in the back halls for three years now?

So I think that's a good comment, and trying to brainstorm about these unfinished thoughts here of having a work party leader actually have sidebar sessions after a meeting, after a work party meeting that can answer questions for people that are just unfamiliar with a topic, or maybe they're new, to bring you into context.

We're generally friendly people, so we would really like to get you up to speed and fill in the gaps to get you to really contribute. Because those valuable opinions and valuable ideas from others that fall through the cracks because you're just unfamiliar maybe with the terminology, the lexicon, the history, you know, that's a shame to lose those opinions. Thanks.

JEFF OSBORN: And we're trying to do that at the end of this session. Abdulkarim.

ABDULKARIM OLOYEDE: Thank you very much. I just wanted to also make a comment that the view of ordinary Internet users who might probably just be using the system and not having any technical knowledge is also very important, especially in what we do.

This is important [inaudible] sometimes when you work on some of these things, you tend to assume some things. So it's also quite important for some of these working groups, for people from the outside who has probably nothing or little to do with the root servers to also give their own opinions. I just wanted to make a comment on that.

JEFF OSBORN: That is a very good idea, and that points to why this is something of a conflict. Because I think we have Bylaws that say a certain thing, and then we have greater needs that say something different. And how to try to serve all of those is beyond my pay grade. I'm not sure how best to do that, and "duly noted" is about the best I can come up with.

Ozan, do we have an agenda?

OZAN SAHIN: Sure. Let me pull up the agenda.

UNIDENTIFIED MALE: [inaudible].

JEFF OSBORN: Who?

UNIDENTIFIED MALE: Wes [inaudible].

JEFF OSBORN: I'm sorry, Wes. Didn't see your hand.

WES HARDAKER: No worries. Wes Hardaker, speaking with no hats. That being said, my many hats actually does give me a wider perspective into: this problem is not specific to just the RSSAC and the RSSAC Caucus.

The SSAC is actually open this week. They opened for the first time, so they've been a closed group for a long time. And I think they have watched—they're the latest of the ICANN groups to become open, and I encourage you to go listen to the SSAC sessions. I wish I could, except that I have too many conflicts that ...

It's another technical body that you can get technical expertise just by listening to. They have sort of a similar problem. They try and measure the activity level of their participants and things like that. I'm getting a large nodding head back from Tara back there, who's now the vice-chair of SSAC.

But it's not just them. Where to participate? If you are interested in participating in ICANN, it's often very hard to figure out how to get started. It's very hard. If you just want to get a increase in your participation, where do you go?

And there's a bunch of groups within ICANN that are—they're open for observation, but you're not actually able to participate unless you run a registry. Right? That's a very wide sentiment. And I think that's one

thing that ICANN can do as a whole, is to give better guidance as where you can come participate.

And actually, the RSSAC Caucus is one of the avenues to be able to participate, as is SSAC, if you're technical. It's less clear if you're into governance or if you're in other things and you're not a representative from a country or you're not running a registry or are not on the business side of the registry. It's very hard to figure out when to get in.

And I will look to ICANN staff members to help me sometime. I think we actually need a bigger list, globally, of: I want to participate in ICANN. And you found the RSSAC Caucus, and you think that's not quite right because you're in that gray bar of not having enough expertise. Can we help redirect you to somewhere else that might fit your skill set better?

And I think this is a bigger ICANN problem that I will look into—will look into with my hat.

JEFF OSBORN: Ozan, we're down to the existing work parties and ongoing work. Did Robert join us?

KEN RENARD: Yes.

JEFF OSBORN: Okay. Robert, if you're there, can we hear about the state of the RSS Security Incident Reporting Work Party?

ROBERT STORY:

I did join last minute, but I think Ken might have actually been prepared for this question, and I'm not. If he's not, I can make something up. But I'll give him a chance to go.

JEFF OSBORN:

Ken's looking pretty game. I think he'll do it.

KEN RENARD:

We'll give you a pass for, what, 2:00 in the morning? All right. No worries.

So the Security Incident Reporting Work Party has been active for two months or so. So this work party came out of the Success Criteria document for a governance structure, which calls for a governance structure to have some facility to do Security Incident Reporting of the RSS.

So this work party was convened to make recommendations to the governance structure, which is still in the future, to make recommendations of what security incident reporting would look for the Root Server System. The main purpose of the work party of security incident reporting is to have transparency for the RSS and give the community an accurate representation of what the security landscape is for the Root Server System.

If you don't hear anything at all about security in the Root Server System, does that mean everything's good, or does that mean that they're hiding something? So we hope to answer that question or at least make it obvious that the answer is: nothing exciting is really happening.

So the work party is looking at things like: what needs to be reported or what should be reported—what types of security incidents? How high is the bar? The answer we're coming up with is pretty obvious that it's subjective.

So we want to document some things about what goes into the decision, maybe just a few examples of what needs to be reported and also look at how should it be reported. If a major incident were to happen, okay, we have a report, maybe even a template to fill out. Some of the good ideas that have been thrown out are periodic reports—monthly, quarterly, semi-annually, whatever period—just to say, "Hey, everything's quiet. Everything's good. Here's a few statistics. Have fun."

But security incident reporting is something that we've seen as an ask from regulators. NIS 2 and CIRCIA are two that are asking for the Root Server System, or even more specifically, RSOs to report on.

So that work party generally meets every two weeks. It did meet here yesterday or the day before. And you're welcome to join the work party. The announcements go out to the Caucus mail list. I encourage you to join. I really encourage you to read the current document before joining to get up to speed. It's not a complete document. It's a rough draft. It's got a lot of comments. It's ugly for now, but it'll at least give you an idea and get you up to speed.

I encourage you to join, encouraged to ask questions on the Caucus mail list. Look at the document. It's not just limited to the work party sessions, the phone calls that are bi-weekly. Please feel free to discuss on the mail list.

And I think, Ozan, you sent out the link to the document with the Caucus reminder?

OZAN SAHIN: [inaudible] putting it in the chat.

KEN RENARD: He's putting in the chat as well, so fantastic. I encourage you to join. And, Robert, if there's anything you'd like to add?

ROBERT STORY: No. I think you covered it pretty well. Thanks.

KEN RENARD: All right. Thanks. Back to you, Jeff.

JEFF OSBORN: For an update on the Document Repository. Hafiz, I'm assuming this is yours.

HAFIZ FAROOZ: Yes, this is mine. I'm actually presenting on behalf of Andrew and Baojun. I don't know whether they are in the meeting or not. Just a quick update about the Zotero platform which was presented in the last IETF. And this is one of the small repositories we built for RSSAC, for the research papers or the related research for DNS/DNSSEC, and other related topics. Next slide. Ozan, if you can help.

This is just a quick [inaudible]. [I recently] joined RSSAC around about one year. So I work in the network and security areas. Next slide, please.

Yeah. For the guys who did not see this presentation before, we started using Zotero as a community research management platform where we will accumulate different research done by RSSAC members or from outside to build a centralized repository. And it was also recommended by Wes in the past that we should also share it with the IRTF. It might be useful for them as well.

So this is a nice industry tool used by researchers for citation and reference management. It's supported on Windows, Linux. [I'm not going to further] details. It has connectors with the web browser if you're using Chrome, Microsoft Edge [inaudible] connector, and it can easily read any research paper metadata and add it to the repository.

And then once your repository is built in, it's nicely accommodated. You can use some plugins for Microsoft Word and Google Docs to start using your repository to speed up your documentation and further explore the existing research work.

And it also supports collaboration through groups, where a team of researchers can share the data. It has cloud functionality, also, and offers around about 300 megs of space. Next, please.

So far we have around about 250-300 research papers. This is [accumulated] by different members of RSSAC. We keep adding different researches we come across, so if you guys have some nice documents to share, you're most welcome.

We are building some tags, nice tags, which will speed up the searching and help you guys if you are working to make a document related to DNS/DNSSEC. Next slide, please.

And it also supports different citation formats—IEEE, AMA—or any other citation formats you're looking for. It's very flexible, a very nice tool for the researchers. Next slide, please.

And this is how you quickly add any footnote you want to add, any bibliography. It's very easy. If you're looking for something, it is there in the repository. So far we are not attaching the PDF, but at least the metadata is quickly available. The URLs, any related conference information, all is there. Next.

So we're, I think, in a very good shape. We started with around about 30-40 papers, and now we are [touching] around about the 300 mark. We will keep adding and keep hearing with the other members.

We encourage you guys, request you guys to, please, whenever you have a document to add, please share with us. Or you can also subscribe to this particular URL. It is a public repository. You can just join in. You can start contributing. And it will be, I think, a nice repository in the future for the Caucus members.

So this is just a quick update. Any points for discussion, most welcome. Thank you.

JEFF OSBORN: That's impressive.

KEN RENARD: Erum. There's a hand.

JEFF OSBORN: And there's a hand.

KEN RENARD: [inaudible].

JEFF OSBORN: Erum.

ERUM WELLING: Thank you. Thank you, Hafiz, for covering that. I have a question. Would it be worthwhile to perhaps include some kind of contact information, email address, so that if there's a research paper that we'd like to get clarification on a particular paragraph or whatever, concept, that here's a way to get ahold of the expert who's provided—who's done the research for perhaps some clarification? Or is that asking too much?

It just seems like if we're getting started on something, this might be a good time to mention, you know, if there's an interest in a particular additional field or metadata that we want to keep. Thank you.

JEFF OSBORN: Wes.

HAFIZ FAROOZ: Yeah, I think it's—sorry. So I think it's a very nice addition. Obviously, we can add some new fields for the email contacts, phone numbers, just in

case if somebody wants to share. Right now we have—even many RSSAC Caucus members, we have some documents from Wes and other Caucus members also.

If they can provide something extra, if they share information, we can also populate that. We will keep in mind, and whenever we get something from the research paper itself, we will try to add it as a field so that it's easily searchable, easily exportable. And, yeah, it's a nice feature. Thank you.

JEFF OSBORN: Thank you. Wes, you had your hand up?

WES HARDAKER: I did, thanks. First off, thanks for putting this together. It's a lot of work to put together a database this large, I know, because we track stuff at ISI, too, for the number of papers that we collaboratively produce. The good news is papers do have email addresses on them, most of the time, not all the time. But a lot of times, [in the header, they're there].

My question actually was going to come centered around how are you handling conflicts as you get upwards into 200? If I gave you 75 papers that my colleague wrote, because he's written a ton of them, could you easily import those and delete the conflicts? Or is this becoming more manually labor-intensive as time goes on?

HAFIZ FAROOZ: Yeah. Zotero is very intelligent. If you have a Firefox plugin, a Chrome blog plugin installed, it reads the meta tags of the website and collects

that information automatically. So we're not actually adding any manual work. Maybe the email address we're talking about, that might be something manual because there's no meta tag, I remember, in these websites. If IEEE Explorer or any forum, if they don't have a meta tag or email address, then we will have to add it manually.

Otherwise, as of now, it's all automatic. Plugin is very sharp. It gets the author's name. It gets the [inaudible] categories. So all this information is automatically imported into the system.

JEFF OSBORN: Very impressive. Any other questions? Andrew.

ANDREW MCCONOCHIE: I think maybe Wes's questions was about deduplication. And I had to use Zotero. I first got introduced to Zotero on a paper I was writing, which ended up having, like, 140 citations. And so I had this problem with duplicates where you get the paper from the conference and then the paper from the journal, and it's the same paper but slightly different at different times.

And what I love about Zotero is that you can just sort by name or sort by whatever, or search and reduce. So the duplicates become really, really obvious. So actually, deduplication is pretty easy because it has all the metadata in it, and it just kind of automatically gets it. But yeah, it's slightly manual, but it's pretty easy.

HAFIZ FAROOZ: Yeah. We do the same thing. After adding it to the repository, we review it. And in case of duplication, we always try to pick the one which has more rich metadata. And, obviously, these are from IEEE Explorer and these websites. They [inaudible] manage the metadata, so that is priority number one. Any other website which does not have rich metadata, we try to remove that duplicate.

JEFF OSBORN: Thank you. Next item on the agenda. RSSAC047. I think Paul was going to join us remotely.

PAUL HOFFMAN: I am.

JEFF OSBORN: Hi, Paul.

PAUL HOFFMAN: All right. Next slide, please. So this is pretty much a one-slide presentation. So a few people were asking what is the status of RSSAC047 since we finished in the initial implementation. We reported to RSSAC at the last ICANN meeting about what it meant for us to complete the implementation but also to keep collecting data.

For those people who have read RSSAC047 and didn't know about this, our complete set of code and instructions are at a GitHub repo. We still are open to getting comments and suggestions. We're not actively

working on code now, but if someone is feeling like working on code, they certainly can go through and do pull requests and such like that.

So RSSAC047, the purpose was to collect data long term and also to report monthly. So we're running this initial implementation to collect monthly reports already.

We don't have the complete complement of vantage points that RSSAC047 calls for, although I covered the reasons for that in the report that I gave at ICANN77. But we have a bunch of vantage points throughout the world collecting data every five minutes, passing them to the central data collector site.

And then I do, every month, creating the report for last month. And as expected, most things pass in the reports with the one exception of RSS availability is still failing. And so I gave an example here from the September report of that. In fact, because we said that the threshold has to be 99.999%, and in the cases shown here, because it was 99.997%, RSS availability is failing. But other than that, everything is passing.

So the next step is that RSSAC can decide what it wants to do about RSSAC047. We did a couple of updates last year, minor updates, but there's still the big question of: are we measuring the right thing? Do we want to be measuring more? That's an RSSAC question.

ICANN will continue to run this indefinitely just so that we have collected enough data to help RSSAC. If RSSAC says, "Where in the report are things failing," things like that. But again, I want to emphasize, this is an initial implementation. If RSSAC wants to do a full

implementation on this, that needs to be a separate contract. Of course, they can start with our code, or they can start fresh.

So that's it for me. I see Wes has his hand up.

WES HARDAKER: Yeah. Thank you, Paul. First off, thank you for presenting at 2:30 a.m. your time. I know what that's like, having done it many times.

One question about the RSS availability that's failing to at least meet the threshold we originally defined. Your example on the screen shows only UDP failing, which is probably not what surprises me. And I'd have to go back and look at the—I won't get into the details of RSSAC047's measurement and whether we should adjust those or not, or whether they're sufficient [in retries] and all that kind of stuff.

But can you tell me, is it only UDP that fails?

PAUL HOFFMAN: Yes. So what I show here is literally the only part of the monthly report that shows a failure.

WES HARDAKER: Okay. So in any of the monthly reports, does TCP ever fail, too?

PAUL HOFFMAN: I haven't been looking that carefully, but I've been sending them to Duane, who's sitting in the room. He might remember better than I do.

WES HARDAKER: Duane just left.

PAUL HOFFMAN: Sorry to scare him away.

WES HARDAKER: But, no, that's fine.

PAUL HOFFMAN: I'm pretty sure that TCP has been at 100% RSS availability. I strongly suspect, without looking at the data—because, again, this is just sort of a maintenance thing right now—that if you tweak to the definition for UDP on timeouts versus things that are known to be routing issues ...

And do remember that part of RSSAC047 is every time we do a five-minute measurement, we also do a traceroute. So if you looked at the few times where the UDP failed and you were really concerned about it, you could actually go back and look at the traceroute for that particular five-minute increment, for that particular vantage point and see. I haven't done it. But, yeah, that still is there.

And if it turns out that this is of concern and people say, "Well, we need something more than a traceroute," those kind of things can be added to RSSAC047. And I'm going to be extreme here. If someone said, "We need to see a routing map right there," that's not something we do now. We could add it later. Things like that.

But at this point, I suspect that the likely thing that will happen is, instead of saying five nines, which is arbitrary anyways, we might actually just change the threshold.

WES HARDAKER: All right, thanks. That's very helpful. There's a lot of interesting data to dive into now. That's awesome just from a pure research and playing with numbers perspective.

PAUL HOFFMAN: Yep.

WES HARDAKER: I wonder if I could this [inaudible].

PAUL HOFFMAN: And so anyone who wants me to send them the monthly reports, send me email, and I'll make a little list. I create the reports by hand, but it literally is: go to the collector box, say "run the monthly report." I have to remember to do that on, like, the 2nd of the month just to make sure all the data is in. And then it is a text file that I copy and paste at this point. I send it just to Duane. Happy to send it to anyone else.

WES HARDAKER: I just loved your reference to "go to the collector, get the box." Thanks for that.

JEFF OSBORN: All right. Ken, is that an old hand up?

KEN RENARD: That's a new one. Is raw data available somewhere? Or can we request that from you if it's reasonable to send around?

PAUL HOFFMAN: Well, that's an interesting question because, in fact, when RSSAC047 was written, there were many people, some of whom are sitting in the room, who said that for the real implementation, the raw data absolutely should not be available. People were worried about the RSSAC047 data being used—I'm sorry, some Root Server Operators were concerned about the RSSAC047 data being used against them and such like that.

So I have not made the raw data available. Again, this is up to RSSAC. If RSSAC feels that because this is an initial implementation, it's okay to do that, that's fine. Until I hear differently, though, I'm going to go with the sense of the room during the last workshop that created RSSAC047 and not make the data available, only make the reports available.

One thing to be clear on, though, is RSSAC047 said the reports must only say "pass" and "fail." They must not give more information. My monthly reports come in two parts: the part where it's only pass and fail, and then also the parts that say, "By how much did you fail?"

So for example, here on the screen, notice that it says—this one's an easy one because it's a percentage. So you could probably count how many data points it would take, how many measurements it would take

to come to 99.997%. But when I distribute the reports, it does have the two sections in it.

But, no, I'm not happy on giving out data without RSSAC itself saying, "Yes, it's okay for us to be sharing data at this point." And I imagine that would be an update to RSSAC047.

WES HARDAKER: Fair enough. Thanks.

JEFF OSBORN: Any other questions for Paul before we ask him to address RSSAC028? All right. Paul, still yours.

PAUL HOFFMAN: Okay. Now this one will be two slides. Next slide, please, Ozan.

So as a reminder, RSSAC028, which we wrote many, many years ago, had some data in it. That is, RSSAC028 had a bunch of proposals for how to change the names of the root servers with the goal being that it would be nice to have the root server name server information be signed.

So Recommendation 2 from RSSAC028 asked the Board to conduct studies to understand the current behavior of the resolvers and how each of the naming schemes would be affected by those behaviors. So OCTO commissioned that study last year. We did not do it in-house. We did it out-of-house. We were very, very happy with the results.

The contractor for it was NLnet Labs, and I'm not going to—there are people in the room whose Dutch is much better than mine, and those are the folks who worked in the contract. So the report that they give is now complete. We sent that link to RSSAC last month.

Actually, Ozan, if you can copy and paste the link from this underlined "report."

OZAN SAHIN: Done.

PAUL HOFFMAN: Very good. Thank you. You will see that the report is incredibly detailed. It has some surprising results. I would say this is not the kind of report where you can say, "Let me just go look at the end and see." The mechanisms that they use were quite detailed.

And I do want to do a shout-out to the root server operators on this. When we did the initial testing for RSSAC028, we just took a couple of open-source resolver software at the time and did it, but we know that some root server operators run their authoritative root server software.

Some of it is proprietary, even the ones who are using non-proprietary software might be, you know, the configuration files they're using might be a little bit odd; or odd by if I were running an authoritative server for a personal domain name, I would have different configuration. So the root server operators were extremely open with the contractor multiple times so that the report could be quite complete.

The report does cover some of the confidential implementations. It does it without naming folks, and it gives them generic letters. But if you are interested in rootservers.net and what might come next, please do go read the report. That is now complete. Next slide, please.

So the next step, actually, is something that we're doing for RZERC, the Root Zone Evolution Review Committee, not for RSSAC. RZERC002. Recommendation 1 in our RZERC002 said, "Go do what RSSAC asked." So that one is now covered.

But the second recommendation is that they wanted ICANN [inaudible]—now that we have a different CEO, we don't have to use the word "Org" anymore—ICANN staff to further explore the cost-benefit trade-offs and risks of the signed root zone name server data. Do the risks of redirected query traffic outweigh the risks of operational complexity?

So essentially, this is saying now that you have a report with lots and lots of data in it, please go do an evaluation on that, and look at the risks. So we're now commissioning this study. Even as we speak, we expect it to be completed by mid-next year. If things work out on the contracting, which I'm hoping that they do, we might have this report completed by DNS-OARC in February and do a presentation there.

Certainly, I would hope that even if we do that, that we would then report back again at the next ICANN meeting in March. So that's our hope. Deadlines don't always work, especially with holidays in between, but we're looking to do that.

And then after that report's done, if RSSAC wants more work being done in the area, that's great. My guess is what will happen is RSSAC will have to then look at the results both of the initial report and of the follow-up report that's talking about risks and also implementation, and decide what they want to do.

This is unlikely to be the RSSAC028 v.2 because RSSAC028 is, "Please look at lots of choices and report on them." And so if RSSAC wants to work on the result, it will probably be in a brand-new RSSAC document, not just a new version of RSSAC028.

And that's it for me on this. Are there questions? Okay. Thank you very much.

JEFF OSBORN: Thanks, Paul. This leads us to the slightly awkward section of the update from Brad who had to leave. Do you want to fill in?

KEN RENARD: I can—

JEFF OSBORN: Do you want to fill in?

KEN RENARD: I'll just say a few words. Maybe we can come back to the new work items shortly here.

But the GWG is the Governance Working Group. It's tasked with making recommendations, building a governance structure to govern the Root

Server System and the RSOs. They are meeting here this week, doing a lot of good work. There's been a lot of progress in developing principles of governance, and we're here this week.

I believe you'll find those are open sessions. Yeah. So those are open sessions. You can find it on the schedule. If you'd like to sit in and listen in, you're welcome to do so. Anybody else want to say anything about the DWG?

Ozan.

OZAN SAHIN: So there's one a.m. session tomorrow, one p.m. session. And then on Thursday, there will be, again, an a.m. and p.m. session. Four sessions here.

KEN RENARD: Yeah. If we could go back to the 4 (a) (v), the new work items. Ozan, can you pull up the—all right. He's got it done before I can even ask the question.

So from the Caucus survey, one of the open questions are: what are some topics that you think the Caucus could look at? So the process within the RSSAC is usually that these topics come in, the Admin Committee does a first pass of filtering, then we bring these topics to the RSSAC which then decides which ones are worth pursuing, and building a Statement of Work and possibly spinning up a work party.

So these were some of the suggestions that were from the Caucus survey. And not all of these really apply to the Caucus because we're

really focused on the Root Server System and how it works, not necessarily DNS aspects that touch on the Root Server System. So we're really looking for topics that would analyze or advance operations at the Root Server System.

So with that said, I just wanted to invite anybody who made these suggestions who actually wants to expand on them further what they mean to do so. If you made these suggestions, you're here in the room or online and would like to say any more about them, this is an opportunity to do so.

I'll just wait for hands. Seeing nothing yet. So we've seen these responses from the survey. We don't know who made them, but we just wanted to give that opportunity.

Wes.

WES HARDAKER: I just want to give my appreciation to those that did submit answers. There are things, as you said, that are out of scope of the Root Server System. But a lot of times even brainstorming comes up with ideas that, you know [inaudible]. So as an example, the Root KSK rollover. That is not our responsibility. We only take the root from the IANA. They're the ones that are thinking about doing KSK rollover studies.

Having said that, there is an element that applies to RSSAC, which is: okay, how's it going to affect us? Well, the study has already outlined packet sizes and stuff. But doing tests or things like that, or even creating a report saying, "Hey, if ISI runs a root testbed ..." So we

actually have the ability to test things ahead of time by messing with stuff.

And so one of the things that we should do, taking it back to a bigger-picture view of this list, even when there's things that don't directly apply to the root, can we tweak the idea a little bit to make it interesting or more relevant to RSSAC itself?

KEN RENARD: Yeah, good point. Thanks. Thank you for those that did suggest these topics. And I think we will keep them around. Whether they can be refined into something that's specific or just used in the future, we appreciate the input.

All right. I think with that, if we go on to the last agenda item. We've got 12 minutes left.

So again, we've talked about the survey. We've talked about people coming on board as newcomers or some that aren't really familiar with the exact topics that we're discussing. Some of them we've been discussing for a long time. We wanted to have a time for Caucus members to ask any specific questions about our work party.

In this case, it's specific to an RSO. We've got several RSOs in the room here. If there are things that you, the DNS experts, maybe are unfamiliar with, with the RSO operations, really anything else in the Caucus, we want to be able to help you to gain the context to be on board and be able to participate more effectively and share your thoughts with us because that's the value of the Caucus, the diversity of opinion.

So I just want to open this up. Ask an RSO, ask an RSSAC member, ask a Caucus member anything that can help you more effectively participate in the Caucus.

Go ahead.

NAVEED BIN RAIS: I have a generic kind of question. Have we ever discussed would there be a time that comes where you believe that the number of instances that the root servers have across the globe are going to be overwhelmed or are more than enough now, and we don't need to place another instance; or it depends on many other factors like the willingness of a particular country to host more than one or all of the instances of all root servers?

So is there any technical kind of barrier to this or a debate on this within the RSSAC?

KEN RENARD: I'll defer, I think, to Wes.

WES HARDAKER: I was going to say, how are we handling questions in the queue versus answers in the queue? I wasn't sure if Erum was planning on answering, or Paul for that—

ERUM WELLING: [inaudible].

WES HARDAKER: No. And I suspect, Paul, you have a question, not an answer?

PAUL HOFFMAN: No, I had an answer, quite frankly.

WES HARDAKER: You're ahead of me. Go for it.

PAUL HOFFMAN: Okay. So thank you, Naveed, for the question. I gave a presentation just a few minutes ago on RSSAC047. RSSAC047 was very much around—I'm sorry. The results that we will get by looking at RSSAC047 was very much around answering your question. Because we have not finalized RSSAC047, these measurements are not being done officially and such like that, but one of the geneses of RSSAC047 is to answer the question: are we doing it sufficiently?

And as a review for the people in the room who haven't read RSSAC047, the report has basically two parts. One which is: is this root server operator meeting the minimum goals that we have set? But the second is: is the RSS as a whole meeting the minimum goals that we set? So there are different thresholds as such.

So to answer your question more in-depth, I would say look at RSSAC047. See whether the kind of information that it would be collecting would be helpful. And if not, think about additional metrics or additional kinds of measurements that we can take.

There is also RSSAC002, which is a self-reporting mechanism which doesn't go to the level of detail that you were just asking, Naveed, about number of instances. RSSAC002 really is about each individual Root Server Operator. And it doesn't even require a collection of all them, although people like Duane have done that. But RSSAC047 is exactly about that question. Hopefully, that's helpful.

NAVEED BIN RAIS: Thank you very much. Yeah.

WES HARDAKER: All right. Yeah, it's a good question, Naveed. One thing RSOs do is evaluate everything based on a technical need. Right? We're a technical bunch of geeks that do things on a technical need. So when we're placing new instances or deciding where to put stuff, it's based on finding some solution for that particular location.

So it's better latency for that area. Sometimes it's protecting against DDoS for a particular area, or anything in those. We add new routes to existing instances all the time just to try and make sure that ...

In fact, probably all of us would agree that playing with routes takes up more time than anything else because you are constantly twiddling stuff to make the service better and better.

You used the phrase, "Do you think there could be too many?" I loved that. It made my brain spin for a little bit. Because I don't think that there's really a limit that's really a conceivable limit that would be too many. Right? We don't hurt things by adding more.

That being said, the current Root Server System is so over-provisioned with instances. There are 1700. The world actually doesn't need 1700. That's more than is needed, but we're doing it to solve other things and to really ensure that the long-term resiliency of the root just can't be taken out. It's going to work. Period. No matter what happens in the world.

So we wouldn't stop. Right? I can't think of a reason why we'd stop adding things because we're always solving these little micro problems. We're always trying to improve the system even though we're really at a pretty great spot.

And I'll bring up one final thing, which is even the concept of hyperlocal roots as the ICANN terminology calls it—or RFC 8806 is where that derives from, that, actually, Paul Hoffman's an author on. I have a local root project that I ran out of ISI that allows you to mirror a root off of us to your local network and things like that. That's even adding a bunch more, but they're not global root instances. They're just a root for your local network.

Those types of projects are adding even more and more, and it just distributes that load and responsiveness even better. So I give that as a data point because it shows the number that could be added from that just keeps growing and growing and growing. So I don't think that there's a notion of too many, but I love thinking about it. Could there be too many?

My answer is no, from a reasonable engineering point of view. Could you be crazy? Yes, you could. And that would be silly.

KEN RENARD: Thanks, Wes. We have Erum and then an online question. Erum.

ERUM WELLING: Thank you. I won't take too much time. Just curious about SSAC. We have a working party, for example, related to security. I wasn't sure what purpose, or if there's an intersection with SSAC whenever we talk about security-type issues. Thank you.

KEN RENARD: Yeah. For the security incident reporting, that's something we've always discussed or brought up with SSAC in our joint meetings. We've discussed having that work party for a long time. We discussed that we were starting it, and now we actually did. So to that end, we've actually really hard-pressed SSAC to have members join the RSSAC Caucus. And there's no question of expertise there. So, thanks. Good question.

                      Ozan, do you want to do the online question?

OZAN SAHIN: Sure. The online question comes from Wataru Ohgai. And the question is, "Any future plan to integrate RSSAC047 and OCTO ITHI Project?"

KEN RENARD: Paul.

WES HARDAKER: I would like to defer that one to Paul, if you're available.

EN

PAUL HOFFMAN: Sorry, I was trying to raise my hand to answer that, but you all asked. There is not. That is that RSSAC047, the data that we have now is quite informal. RSSAC might come to us and say, "Hey, can you change all of that?" They might say, "We were wrong. We want weekly reports and such." So again, what we've done is an initial implementation for the benefit of RSSAC, and we're waiting for more guidance from RSSAC to tell us what more they would like and things like that.

ITHI is really about looking at current data around the world from different sources that is meant to be much longer-term. Nothing in ITHI would be used as an alert, but the purpose of ITHI is someone saying, "Well, wait. This number seems weird. Tell me about the last five years of it." So they really aren't compatible that way.

But, for example, if RSSAC said, "Great. Let's really instantiate RSSAC0147. And how are we going to handle the reports," obviously, ICANN staff would help with that. We would be able to set something up. But they really are distinct from ITHI. Thank you.

JEFF OSBORN: Joe.

JOSEPH HAYES: Hi. This question is for those who feel that they don't have the technical skills needed to help with RSSAC Caucus or just want to learn more about DNS in itself. So Red Hat is available via developer license, and

78 ANNUAL GENERAL MEETING

BIND is open source. What can someone do at home to teach themselves DNS infrastructure?

KEN RENARD: I'll take a stab. I can wave the ICANN flag. There are some actual ICANN Learn courses that you can take. There's some very basic ones on DNS itself. I don't think they'll get you to the typical technical level of the existing Caucus members, but it's certainly something. My advice is: play in a safe environment.

Wes.

WES HARDAKER: Yeah. It's a great question, and we've been thinking a lot about it lately. How do we bring about more education to the system as a whole? And so there's a few things that I can probably point out.

One, there's actually a number of good YouTube videos and things like that, that describe the DNS. And interestingly enough, not too many concentrate so much on the Root Server System.

But one thing that you can do is you can actually go look at—even the RSSAC history document has, actually, a lot of buried little technical details and stuff like that.

The other thing that I mentioned earlier is running a hyperlocal root yourself. If you really want to know more about the root, that's the best way. Right? You're functionally running a copy of the root zone in your house or whatever. I have one running in my house because I'm a geek, and I do those kind of things.

But the other thing is, you know, there's ways to sort of, if you really want to get super detailed, you can go run your own copy from scratch. There's ways to mirror the data from other places just to go tinker with it. Obviously, you can't advertise BGP routes to existing addresses to the world. That would be bad. But you can actually build toy networks all the time.

Or ISI has a research testbed, too, that we offer to let people come use. And it's actually a mirrored copy of our real production system with the same hardware and everything. You have to have a research purpose to come in and use that. But there's lots of resources out there.

You're welcome to contact me offline. We're a university, so we're always here to help people. So do contact me.

JOSEPH HAYES: Like I said, this is on behalf of the people out there who want to learn more information. But, thank you. That's exactly what I was hoping somebody would say.

WES HARDAKER: Yeah. Anybody's welcome to contact me and I can help where I can.

KEN RENARD: We're already one minute over. We have Paul Hoffman online, and then one in the room.

EN

PAUL HOFFMAN: Yeah. So I recognize that we're over. But very briefly, to answer the question of: are there resources for helping people learn how to set up an authoritative server, not just a YouTube video but hands-on?

A different part of OCTO, the technical engagement part, does this regularly. We will go, generally, to a developing country or developing region, maybe to a NOG. And we have an entire setup where everyone sits down and starts hitting the keyboard and setting up authoritative servers and trying out stuff.

So that's not generally available now, but if there is a strong desire for that, we could probably operationalize that out of it. It certainly does take some hand-holding. But we, on a very regular basis, will be training dozens and dozens of ISPs and such like that in regions about: what is the DNS? How do you set up authoritative servers? How do you set up recursive resolvers—with a hands-on setup.

KEN RENARD: Thanks, Paul. Go ahead.

YAZID AKANHO: Thanks, Ken. And thanks, Paul. This is Yazid speaking from Technical Engagement. It's exactly what I wanted to also say.

PAUL HOFFMAN: Sorry. Yazid, I didn't know you were in the room. I would have just pointed at you. Please. By the way, for the folks in the room, Yazid's the one who set it up.

78 | ANNUAL GENERAL MEETING

YAZID AKANHO: Thanks, Paul. Thanks, Paul. So we're actually a team which is called the Technical Engagement, of course, a part of the Office of the CTO. And that's precisely our work.

I'm right now setting up the environment to deliver remote hands-on tomorrow afternoon for some communities in Southern and Eastern Africa region. Since three years, we have been doing that. Even during COVID, all our technical engagement were remote, and it went well.

We have also supported a lot of ISPs in configuring and turning on DNSSEC validation, for example, and also supported even some registry operators on signing their ccTLDs; and definitely NOGs and other kinds of universities as well in understanding the basics of the DNS deployment.

So we are more than welcome to support the RSSAC and the RSSAC Caucus if [you need this]. Thank you.

JEFF OSBORN: Thank you very much for that. I thought this was going to be a short meeting. It has actually run long, so I appreciate all your time. Oops.

ANDREW MCCONOCHIE: One AOB. If you haven't yet signed the attendance sheet and you're in the room, please sign it.

JEFF OSBORN: Going once, going twice. Thank you all for attending.

OZAN SAHIN: Please stop the recording.

[END OF TRANSCRIPTION]