ICANN77 | PF – RSSAC Work Session
Thursday, June 15 2023 – 13:45 to 15:00 DCA

OZAN SAHIN: Hello and welcome to the RSSAC work session on the security incident reporting. My name is Ozan and I will be the remote participation manager for this session. Please note that this session is being recorded and is governed by the ICANN expected standard of behavior.

During the session, questions or comments submitted in chat will only be read aloud if put in the proper form. as noted in chat. I will read questions and comments allowed during the time set by the chair or moderator of the session. If you would like to ask your question or make your comment 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.

This session includes automated real time transcription. Please note that this transcript is not official or authoritative. To read the 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 into 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.

With that, I will hand the floor over to Ken Renard.

KEN RENARD: Thanks, Ozan. Welcome. I apologize. I keep getting kicked off the Zoom session. Problems on my end, so if somebody could help me with raised hands, throw something at me, whatever, I will get to you. So welcome. This is the last session for RSSAC for this meeting, and it'll probably be a short one given the size of the audience in the room here. The topic is security incident reporting.

And if we can go to the next slide here. RSSAC has been talking about security incident reporting for a while. We put together a statement of work for an RSSAC caucus work party. That statement of work was voted on and approved just a few hours ago, and it's certainly not enough time to officially convene a caucus meeting. So this is a kind of a placeholder. It's more of an RSSAC session. I do believe there are some caucus members online that were invited kind of last minutes. But note that this is not officially a work party meeting just because we haven't had the advance notice to get everyone involved.

So the purpose here is to try to level set to the RSSAC caucus and to observers about security incident reporting work and this work party. Like I said, we've been talking about this for a while now. We developed a statement of work and that's out there. It's been approved. I believe Ozan put that in the chat for the session. But as we were developing the statement of work , we kind of cracked open Pandora's Box and started to dig into this a little bit for the purpose of trying to put guardrails and make sure the statement of work was where we wanted to be.

So we've discussed this, we've discussed some of the few of the details here and there ad hoc, and we wanted to share some of that, level-set, and hopefully give this session as an introduction to work party members to get them bootstrapped, and ideally (I will not promise) get a document out for caucus members that want to join the work party that they can just read and get up to speed with what has been discussed already, not as rules, not as direction that the work party should go, just as historical … what we've talked about previously. I don't think that this is going to be a real easy work party. I hope to be wrong about that.

Just a reminder that this recording will be publicly available and there's going to be a wide variety of listeners, whether they're caucus party, work party, members or not.

We can go to the next slide. Again, establishing a baseline for what the work party has done, I'd like to go around the room to those RSSAC members that are here that have discussed this previously, to ask you about specific points that we've discussed before that you want to be a part of this baseline for work party members; maybe topics of security, incident reporting to address topics to avoid constraints, maybe even what you think is the success criteria. What would the successful work party outcome look like? Again, it's just to develop something going into the work party. I do strongly recommend that RSSAC members and specifically RSOs actively participate in the work party. While the work party will develop recommendations only, this could turn into things that could be requests or requirements of fruits of our operators. So we want to get this right. Yeah, so we've said this before.

If we go to the next slide, Security Incident Reporting makes a nice acronym, SIR. So again, we start with RSSAC058, which was the success criteria for the RSS governance structure. And you can read that there: "RSS governance structure may include provision for cyber, incident oversight and disclosure obligations," et cetera. In parallel, government entities can have an ongoing and legitimate interest in this. What I read from that is that governments are recognizing that DNS and the root server system are critical pieces of Internet infrastructure. They're looking out for their best interests of their constituents, and security stability, and reliability of the RSS is important. We feel the same way. Our constituents are the entire Internet. We care very much about this. So security incident reporting is a way to be more transparent from the RSS. We on the inside kind of know what's going on. We're confident. We sleep at night. How do we convey that to the rest of the world to maintain and develop that trust?

Go to the next slide. So motivation here, again, is to establish and maintain trust. We would hope that there's a lot of trust already in the root server system, but there are communities that are maybe unfamiliar with the RSS or have never given it a thought before that may have a new interest. So we are establishing that trust with those that are new to this community. So, yeah, again, we know how resilient this system is, but as a security person, with the security mantra if "trust but verify," this is one way to give people the ability to verify. So I want to approach this, with what is the right thing to do for the RSS in transparency, in developing trust.? And that's kind of the motivation that I see.

If we can go to the next slide. So this is kind of getting into some of the topics that we've discussed thus far. Some of these are in the statement of work . Some of them are not. RSOs … We have been reporting, in quotes, security incidents, albeit informally since the inception of the route. So there are public statements. There's a link to an example of one. So very little of the events that have happened in the root server system would rise to the importance of public disclosure. But that's kind of our own definition of importance. So let's at least formalize that. And what level of importance warrants disclosure? If we have a very well- defined threshold, if we say nothing, then you could assume that everything's below that threshold of severity. So yeah, what does it mean if we're silent on security incidents?

One of the things we pointed out is IMRS, the ICANN-Managed Root Server, does something like this. It's a reference. In this slide, it says … I need a reference. I can't point to that document off the top of my head. But they put out essentially monthly documentation of significant events with their root server. Hey, that's a point, a data point, that we could use. As RSOs, as RSSAC, we have significant monitoring infrastructures. There are external monitoring infrastructures as well, such as RIPE, such as Paul's RSSAC047 work . So there's a lot of eyes on the root server system. And all of these are available to us, available to the community for transparency as well, and seeing how well the root server system is doing.

All right, we can go to the next slide. So these are some of the things that have come up along the way. And I'll just go through these briefly, again, giving people the baseline. This is what has been discussed, again, not

as a not as a rule, not as “These must be followed,” but just these are the kind of baseline where we are, where we've come from here.

Obviously, the security incident reporting shouldn't affect the operations of the system. The availability and the service that we provide is very important. We realize that, and we should not have an effect on the actual operations.

And then there's this balance of providing sufficient information to the community, to the world, to maintain trust, but not disclosing too much information about security incidents that would then increase security threats. So there's that balance.

Independence and autonomy of RSOs is documented. We can put in references for that. So we know that RSO organizational policies will come into play here.

So one of the ways we've talked about what makes a reportable incident is something that has a material effect or material adverse effect on the root server system. I purposely put in “define” and scratched it out. Let's put “characterize” because we cannot define “material effect.” It's way too complicated. But we can probably characterize it. What does it look like? What does it smell like? So this work party can hopefully make some headway on that. And it's not defining very specific details but general ideas of what should and should not be reported.

And probably the biggest thing here is that we focus on the root server system itself versus individual RSOs. And DNS users use the system, not the RSOs. It’s semantic, but it's important there. And really there are

security resiliency measures that are in place at the system level in addition to at the RSO level. So make sure we're measuring the right thing, make sure reporting on the right thing.

Can we characterize our audience? Who do these incident reports go to? It's in theory, the entire world, but is there any specific audiences that we should gear the language towards, gear the formats or anything specific to? Okay. And that goes into what type of information would be good for the audience as well. Think of something very simplistic like traffic light: red, green, yellow. Everything's good, et cetera.

So those are things we have discussed so far.

Next slide. One approach that I've had on how to approach this … There's lots of different things. Security is a very wide topic. But going at this through the kind of traditional confidentiality, integrity, and availability approach to security, that was my thought on at least breaking this down into how we can compartmentalize things and think of it in manageable chunks as this is just input to the work party. This approach could be picked up and run with or it could find a better way. That's fine, but we're just documenting what we've looked at so far.

Next slide. Another important piece of this is, what roles should the governance structure have? Merely an aggregation point for disclosure from RSOs? Or basically operating a website where people can see the current security status? Should it be a facilitator of inter-RSO discussions? There are things that we, as RSOs already talk about, already have a very good communication channel for. Should the governance structure essentially take that over or leave it as is? It seems to be operating pretty well amongst us.

And then the last slide here. I'd like to open up the floor to the folks here at the table. Are there any new issues that we've thought of since the statement of work was finalized or formalized? Are there any other seedlings of topics to pose to the work party? Everybody here is welcome to join the work party. If you're part of the caucus, that's great, but this is a chance to kind of fill in … This was my recollection. What else is what else did you have that should be given to the work party to bootstrap?

UNIDENTIFIED MALE: Hand from Erum.

KEN RENARD: Erum, yes?

ERUM WELLING: Thank you. So with incident reporting, a lot of times when we have that in our individual lives, there's requirements that are coming from some external source that's saying, "Thou shalt do this here." Here, it sounds like it's perhaps more of an open book. So is there a concept that we're going to be exploring as far as the audience, perhaps multiple different audiences, for this incident reporting? Because if we report at a certain level, it's not going to make sense to a particular audience. So is this going to be customized for different audiences? Is that kind of the goal—the transparency? And with that transparency comes to different audiences? And then we would embark on coming up with the right reporting sets based on the audience? Thanks.

KEN RENARD: Yes, I think so. Part of the work party is to define an audience. My initial thoughts are that this is for the Internet, for the good of users of DNS. So I don't think we're going to get down into 52 different audiences. It's going to be very broad swaths, but that's for the work party to determine what those audiences are.

I think Robert was next.

ROB CAROLINA: Yes, thanks. A couple of comments, and I'm probably repeating one of the things that I've said before. And from the slides that you've just shown, I would urge the work party that's put together to consider two different frames of reference with respect to the problem of reporting incident reporting. The first frame of reference is, what should the "to be created" RSS governance structure disclose to the world about incidents that have some kind of impact on the RSS?

And then I think the second frame of reference, which would be easier to address after that first one is addressed, is what should individual RSOs be disclosing to the newly created governance structure so that the newly created governance structure can take a view on risk management and try to avoid and limit the risks of having incidents that the governance structure, in turn, would need to report to the world?

One of the reasons I'm suggesting those two frames of reference is that I think it's accurate to say that success criteria 8.1.1.1 in RSSAC058 success criteria was originally developed in response to the emerging threat of regulatory intervention, which could theoretically impose reporting obligations on all twelve root server operators and the greater

risk that the twelve root server operators would in turn be subjected to two or ten or 100 or 150 different sets of reporting obligations. I think there was a hope that the new governance structure could form a clearinghouse for disclosures that had an impact.

So that's one thing. Those two frames of reference, I think, should constantly be kept in mind because I think they're both important. They're both important in different ways.

The second thing is that I'm not convinced that the trigger for reporting incidents related to the RSS as a whole should be incidents that have a material adverse effect on the root server system. I mean, to use the wrong metaphor yet again, shouldn't this governance structure also be reporting near misses? The reason I ask that question is that in past discussions when I ask questions like, “Has there ever been an incident that had a material adverse effect on the root server system as a whole?” usually the answer is, “No. We've had some near misses. We've had some things we had to react to and respond to very quickly,” particularly the DDoS attack that you all are familiar with that I'm not.

So I think that the trigger for public reporting (and in my mind, I'm thinking of it as a public reporting requirement) should be set for the whole of the RSS at some level lower than “has a material adverse effect on the service[,]” but something “to be defined.” That means, “if left untreated, in the absence of intervention,” or, “This is the kind of thing that, if left unchecked, could have developed into a problem.”

So I think, for the public reporting threshold, there's an argument that it should be held quite low. I mean, if we're going to hold up to the world

that we've got a [nine-nine's] availability record, then the goal becomes, how do you … Yeah, I'm done.

KEN RENARD: Right. For example, to date (we're all clapping for you), the referenced disclosure that's on the root ops website did not have a material effect on the root server system, yet it was reported by, "Hey, look, we tracked this, we know this, we watched it. Have a nice day." Those things I think, are absolutely possible if the work party wants to pursue that. Certainly the material adverse effect should, shall, must, (whatever; pick your verb) be reported. But anything else? Like you said, that lower bar is informational. It can absolutely go towards improving trust in the system. So, thanks, Robert.

Brad and then Hafiz.

BRAD VERD: Thanks, Ken. I want to expand a little bit on Robert's language, which I agree with wholeheartedly. Be honest, I've been kind of concerned about this one since the beginning. I know we voted on it and agreed to it, but I'm still a little worried.

My train of thought on this is a little bit different in the sense that I think we need to maybe expand the scope to decrease the scope, if that makes any sense. And what I mean by that is I feel like the output of this work party … I know Robert said there's two approaches, and I agree-- or two Facets. I think there's maybe a third. And I think the third is me increasing the scope to maybe limit the second two pieces. And this goes to RSSAC047. I feel like there should be a recommendation. I'm not

saying this work party should do the work or should create the call it a dashboard or whatever you want to call it. But I think the way you need to start in building trust in the root server system, besides the fact that this black box is just over there and has worked for a long, long time, is to describe business as usual. You got to show metrics. Whatever business as usual looks like, you got to show graphs, you got to show queries, and whatever those stats are (not "whatever," but all the stats that we're collecting in 047). And I know Andrew put together a nice dashboard and whatnot, but I feel like there needs to be some center point that people can look at to understand what business as usual looks like for the RSS.

And then once you start kind of showing people what business as usual is, you go into these other two facets, which is … I agree. I don't think it's material impact on the route that should be reported on. I think it's precedent-setting events, things that are outside of business as usual, something that happens. And that doesn't mean … And what I'm really concerned about in what this group has done numerous times, every time we've had this discussion, is we seem to get wrapped around the axle of "Well, is it a DDoS attack? And how big is the DDoS attack? And what are the metrics? What do we have to report on?" Maybe that's down the road a long, long way from now, but I think you first got to kind of show what business as usual looks like for the RSS.

And then there's a subjective call as to when things happen. The RSOs, like they have so far, need to report on it, and maybe they need to do some more of that, I don't know. And then if there is nothing happening, say that: "Nothing worth reporting."

And then going towards the RSS GWG, I think I would focus on that as a clearinghouse. Again, I think that's a long way down the road right now. I wouldn't rely on that to answer any questions for you in the near future. And I guess what I'm saying there is, I would leverage whatever the RSOs have in place right now to get that done and get that moving versus saying, "Well, that'll be addressed when the GWG gets their work done."

KEN RENARD: Thanks, Brad.

Hafiz?

HAFIZ FAROOQ: I have a few comments on the discussion based on my cybersecurity experience, I think a bigger organization like ICANN generally has multiple incident response plans. One, we have a high-level incident response plan at the ICANN level for example. And then we have a small-level incident response plan which is having a limited scope like RSSAC. So sometime bigger organizations … Like, I'm coming from Aramco. My organization does similar stuff. We have a high-level incident response plan made by the company. And then smaller business lines have their own Incident response.

So similarly, when we talk about RSSAC, RSSAC, first of all, needs to identify the risk. And I think the risk is the root service itself. So whatever we are going to put in the incident response plan, this will be around only about the root service itself. The scope has to be well defined. This is point number one.

And the second thing about escalation, which Rob was also highlighting, is that where the information should be shared, whether it should be escalated to the management or it should be escalated to the RSOs. Honestly speaking, the guideline should come from your high-level incident response plan. And if it is not there, then we can build our own escalation metrics based on the classification of incidents we are going to have. Like, if there is a critical incident, maybe we will escalate it to the management. If this is a minor incident, we may escalate it within RSO itself or within RSSAC. So this is my comment about the escalation and about the scope of the incident response plan. Thank you.

KEN RENARD: Thanks. My take on this is we're really focusing on just the reporting side, what goes out to the world. We're not defining the incident response plan. Those things are well defined within the RSOs and it's really just the transparency of the current state of the root server system and getting that out to the public.

I think Wes was next. Or David. Go ahead.

DAVID LAWRENCE: Yeah, I think that I probably agree with Brad. And I guess my view is that it'd be best to try to establish and publish a baseline of activity because I think one of the goals here is to provide reassurance to various nontechnical bodies that everything is operating nominally, that nothing bad is happening. The question about material effect gets complicated because you have to define to some level what that

actually means. I mean, there have been incidents that have been noticed at the root level that have been noticed by external parties. The ones that I'm thinking of off the top of my head is where various root instances didn't sync up and were advertising or providing data from expired zones (or not yet expired, but old zones) that did result in questions being made public about, why is the data inconsistent? I don't think there's been any denial of service that had an image that anyone would actually notice unless they were actually monitoring the root servers. But given the proliferation of IoT devices that you breathe on and they fall over, it wouldn't surprise me that there's more fun in that future.

So I think the best course of action in my mind would be to be as public and open as possible modulo the concerns of providing information to attackers in the information that's published, just to allow the RSS GS, ICANN, and ICANN's allies to be able to point to that data and say, “Look, this is operating okay. These aren't the Droids you're looking for. Go think about blockchain,” or something like that.

KEN RENARD: A question to the group and to the work party eventually (something a document and put forward) is, as you mentioned data being out of sync, is that a security issue? Is that an operational issue? Let’s discuss in the work party.

DAVID LAWRENCE: Sure. The one case that I'm thinking of was a result of a DS record not getting updated for a particular top-level domain. And one could argue that if anyone cared about DNS, that could be a security issue.

KEN RENARD: Thanks.

Wes?

WES HARDAKER: Thanks. I agree with all of my predecessors pretty much. It's been a good conversation. One of the things I realized is we keep talking about regulatory bodies coming and imposing regulations on us. And I think the one thing we have to remember, because we keep saying a regulation body or somebody's going to impose something on us, is we are the regulatory body for the RSS to some extent. It's not quite that cut and dry, but the reality is our goal should be we're imposing things that we want to put on the root server system as policy on behalf of the world at large. So we are (ICANN specifically in a larger scope) is the sort of body to do that.

Now, the tricky part, of course, is where to draw the line for reporting. That's been [hinted at], and that's what we definitely have to dive into.

One thing that I like to sink into people's minds is, at USC ISI's root, we actually have a line where we report, and we report at different levels. We report to Root Ops first if it's small. We have discussions all the time. It's like, “I'm saying this. Are you guys saying this?” And we do that all the time. And then we have a bigger reporting threshold of it goes on

our newsfeed. And the important thing about that is that, if you're going to make those lines, you have to have the ability to use metrics and measurements in order to define where that line is.

So certainly we're not going to do that today, but there's different classes of events, as has been discussed with my colleagues to the right, especially. And how do we do that? Fortunately, we actually have a lot of measurement related stuff that's happening these days, both with Paul's new work … And DNSMON has existed for years. And there is a benefit to saying multiple levels. "We saw something happen. You may have noticed it because you might be measuring similar things, but nobody was really affected," right? Maybe one resolver failed to get one query he answered and that's it. That's sort of all that's typically happened in the past, whereas with a hard failover, I think everybody agrees, if the root actually went down, it'll never happen. Based on our current architecture, it'll just never happen. But if it did, that's the easy one. And if we're going to define that as the line to write something up, then we're never going to write anything. This whole exercise is pointless. But it is important, as others have said, to say, "Here is an event that happened. There was no visible impact, but we're reporting out of the due diligence." That's what a lot of incident reports do. ISI actually … We publish data for significant attacks so that people can actually download the data and do analysis because we want to promote the future expansion. So not only do we report an event, but we report "Here's the data so you can go look at and analyze it and give us feedback if you want or use it for your own research." But this is FYI.

KEN RENARD: Thanks, Wes. Yeah, I think a lot of that is actually encompassed in what IMRS does. So that's, again, a good baseline.

Paul?

PAUL HOFFMAN: So this is Paul Hoffman not speaking for IMRS at all, because I'm not part of that organization, but I am speaking as somebody who is well known for trying to get good definitions both in the DNS and such like that. And I'm going to be really critical here, but the definitions in the statement of work and in the slides have wandered all the hell over the place. It starts off with security, then basically all we talked about was availability. But it's not really availability because we're sort of saying it's stability and such like that. You can have security incident reporting, but you can't do that against a baseline. The baselines that we have, like in 047, are availability. We don't have a baseline for security or how much have we been hit, unless we start doing a lot of per-RSO reporting. And yet in here it's like, oh, no, this isn't about the RSOs, it's about the RSS.

And so I don't think you have the definitions down. Well, I love the idea of doing baseline reporting for availability. It's why I spent a lot of time on 047 or why I kept spending time on it after I was tired of it. But that's not relevant to security incident reporting. The only time it could be relevant is if the availability of the root server system was being impacted significantly. And we don't see that that happens because we've built it super resiliently.

So I propose, before you start talking about the kinds of things you want to collect, you be really clear if you're talking about security or availability or stability. Specifically, the thing that David brought up was a stability issue. If you don't do that, I think you're going to have reports that are easy to read but impossible to interpret. And I don't think that that's what you want if you're going to be putting a lot of work in. Thank you.

KEN RENARD: Thanks, Paul. Yeah. When you think of the root server system, availability is the first piece that comes to mind. But we just should not ignore the others. They may be mostly no ops, but we should as the work party look at them and try to define what could be reportable; what classes or characterization, not going down the rabbit hole of every particular type of incident.

KEN RENARD: Brad?

BRAD VERD: Thanks. I'm going to push back on Paul a little bit. As I stated, I thought I was concerned with this from day one. I think that the scope should be increased and then the scope should be decreased. And I stated that, on creating a baseline, the main reason for doing that is to begin to build trust from the community, from whoever you want to define as the community. Call it your 52 audiences or whatever. I've heard a number of different … Like, who is the audience? The audience is to build trust. That's it, right? And right now it's a black box. Everybody's

like, what's going on? And there's been a lot of hand waving. So you've got to create the baseline so people who are not in the know can begin to kind of get a picture of what's going on.

And then you step into this other aspect of what's reported and what's not reported. And again, I think if you're putting down some calculation that has to be triggered, I think you're setting yourself up for failure out of the gate. That doesn't mean you don't end up there at some point, but I think all of this we have to kind of learn from the baseline as to what business as usual looks like. And in most cases it's going to be, well, it's business as usual. To build that trust, you might want to just over communicate for a while, right? “Hey look, there was a [2GIG] attack. And business as usual. Hey look, there was a [Syn] attack. And business as usual.” And then that kind of becomes like, “Oh, okay, we get it,” right? But right now it's a black box.

And so … I'm trying to think of the right word. I would be concerned, but I would warn, I guess the work party, to not be too prescriptive on that. I think that is something that you need to mature into, and we're not there yet, in my opinion.

PAUL HOFFMAN: And can I respond, since Brad said he was pushing back against me, but he wasn't because I fully agree? I think that starting a baseline of availability is the right way to go, but then not to say that that is a baseline for security. Later on, as Brad said, reports can be added, but they should be carefully labeled as, “Is this an availability issue coming up?” such as we turned off a bunch of our instances, or they got turned off on us. Or now we are adding in security incidents such as a DDoS

attack but while showing that it didn't affect the line, such as a software failure. That didn't affect the line.

So I'm in full agreement, and I don't want to sound like I don't want to start with a baseline. I just want it to be correctly labeled.

BRAD VERD: So my pushback was more around this need for definitions in the statement of work and being crisp and clear in the definitions because I feel if you're crisp and clear in the definitions, you've already kind of, like, stubbed your toe, because I feel we don't know what we don't know. And so all I'm trying to say is that maybe this is a multiple phased project. Maybe the first phase is to say, okay, we need to create a baseline, and we need to get visibility from the RSOs into certain things. And it comes to a group, and they evaluate it, and maybe they report on it, maybe they don't. And so we just kind of see what happens over the course of a year, kind of pull the curtain back a little bit, just a little bit. And then after that, you kind of say, "Okay, now we've got an idea of what business as usual is. We got an idea of certain events that go on. Now maybe we get into a prescriptive mode," but I feel like that's way down the road.

PAUL HOFFMAN: I can agree with that. Yeah, absolutely.

KEN RENARD: Thanks.

Ozan, if you don't mind going to slide two, then we can go to Russ. I just highlighted that I don't expect this work party to be easy, so I think we're demonstrating that right here.

Russ?

RUSS MUNDY: Thanks, Ken. I wanted to just point out that I think we already have one clear example of people using a word, and different people meaning different things. And that word seems to be really important to this discussion. And the word is "security," as I hear Paul talking about security and availability being two different things, and I hear other people just saying, or at least inferring, that security includes availability.

So the basic terminology … I think we have to be careful about our definitions of … And perhaps the most important term to get definition on is itself, security. What do we mean in this context? Thanks.

KEN RENARD: All right, Russ, I'll put you on a spot. Is availability part of security?

RUSS MUNDY: Yes.

KEN RENARD: Okay. That was just a personal opinion.

RUSS MUNDY: Personal opinion on my part, too, but it's the opinion I've held probably since 1980.

KEN RENARD: Yeah. And again, that's the easiest piece to think of with respect to the root server system. [Is it] available. But integrity or correctness is also extremely important. So availability could be the first one that we address because it's easiest to wrap our heads around.

RUSS MUNDY: In my opinion, Ken, the three things are the critical pieces of security for this context, and that's integrity, availability, and the correctness of the data.

KEN RENARD: All right, so this has been some good thoughts spewing, and this is exactly what I was hoping to get. So this work party will kick off at some point. We go start turning the crank from the SOW into a formal work party. I would encourage all RSSAC Caucus members and RSSAC members to join. It'll be fun. I see Paul's face. That's my view of that this will be fun. It's going to be tough. It's going to be tough. And if we split it into manageable pieces, that's fine. We've had work parties that can't get to a conclusion, and I hope that's not the case, but I think this is something that would be good to demonstrate. Open that black box, be transparent, and demonstrate that the root server system is trustworthy. It's secure, it's resilient, and this is one way of getting that out there.

Any last thoughts?

Next step would be to—go ahead, Andrew.

ANDREW MCCONACHIE: I was just going to suggest, if we're done talking about content, is it that time to talk about Doodles and meeting cadence and that kind of stuff?

KEN RENARD: We usually do that in the first meeting of the work party. I think that would be fine to do that. I think the next thing would be to schedule the first work party meeting, make sure, on that statement of work, now that it's been approved, to get it out to the caucus mail list. And is that good to schedule the work party meeting soonish?

All right, well, thank you, everyone, for the thoughts, and I'm glad we could end a little bit early. If this is your last session, great. Have a safe home, and we'll see you in the work party. Thanks.

ERUM WELLING: Thank you so much.

WES HARDAKER: And thank you so much to all the ICANN staff that made these wonderful meetings perfect.

[END OF TRANSCRIPTION]