ICANN74 | Policy Forum – RSSAC Work Session (1 of 2)
Tuesday, June 14, 2022 – 15:00 to 16:00 AMS

OZAN SAHIN: Hello, and welcome to RSSAC Work Session 1: Cyber Incident Oversight. My name is Ozan Sahin, and I am 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 the session, questions or comments submitted in chat will be read aloud if put in the proper form as noted in the chat.

If you would like to speak during the session, please raise your hand in Zoom. When called upon, virtual participants will unmute in Zoom. On-site participants will use the physical microphone to speak and should leave their Zoom microphone disconnected. For the benefit of other participants, please state your name for the record and speak at a reasonable pace. You may access all available features for the session in the Zoom toolbar.

We have an overflow room called Kilimanjaro located across this room. If this meeting room reaches its full capacity, ushers will have additional in-room participants to the overflow room. With that, I will hand the floor over to Ken Renard.

KEN RENARD: Thanks, Ozan. Welcome to the first RSSAC work session. Today's topic is Cyber Incident Oversight and Disclosure Obligations. So the idea here of the agenda—next slide—is to talk about brief introduction of the topic, what it is where it came from. And the goal of the session really is to decide what should RSSAC do with this? Do we want to do anything, make some advice, etc.? I think to do that, we can go in and actually start talking about some of the issues involved in cyber incident oversight, what they might mean. It's meant to be a very open discussion with everyone and deflect some ideas, and then come back to the topic of really what do we want to do as RSSAC? If we decide to do something as RSSAC, the discussions that happened here today can be captured and fed to wherever that goes. So if we go to the next two slides ahead.

So this topic came—it's part of RSSAC058. The success criteria states there, "A.1.1.1: The Root Server System Governance Structure must include, etc." I'll let you read that. It's important to note that RSOs have been doing this informally since the inception of the root and this is not something that's completely new to us. It's a matter of formalizing it as we formalize the governance of the root server system. This topic was introduced into RSSAC058 around the time that NIS2 was being discussed the proposed incident reporting requirements for DNS operators, which it did include the RSS at that time. We

speculated that this would be one of many regulatory bodies that might require such incident reporting and got us thinking of what are the possibilities here, what might happen. Next slide.

So in this vein, as RSSAC, we've spun up this governance structure discussion. After publishing RSSAC037, we've published additional documents that are advice or additional input to the development of a governance structure for the RSS. For metrics, it's 047. For definition of Rogue was in 056. Is there advice and recommendations? So the advice of RSSAC to give to the GWG or the governance structure to let them know what we're thinking about. This topic of cyber incident oversight has not been discussed in in-depth yet within the RSSAC or GWG. But being somewhat new to the success criteria, we know that some formal requirements for cyber incident oversight will be defined at some point. So thinking about this topic at a very high level, the incident oversight disclosure is a good thing for transparency, further establishing trust in the root server system. But we need to balance that transparency with autonomy and security. And the individual organizations and operational security, we don't want to compromise that as well. So next slide.

Again, our goal is to ultimately try to decide what we want to do with this topic. Here's some possible actions. We could literally do nothing. I don't know how good the cookies are outside, if that motivates people to do nothing. One logical option would

be to spin up a Caucus Work Party to develop some recommendations or advice. Given the security nature of this, we could even invite SSAC members or defer this to other groups such as Root Ops or the RSO consultation groups or any other suggestions here. This is meant as a discussion, meant as let's figure this out what we want to do, and pros and cons of that.

So in reality here, each RSO is going to have to go back to their own organization, certainly with disclosure and security requirements. There are some limitations that each organization will have, and those need to be worked out. So that's going to be part of this process as well. We go to the next slide. We just look and see if anybody has any initial thoughts what we should do with this before we start diving into some of the specifics of it. Wes?

WES HARDAKER: Thanks, Ken. I do think that that's an issue that it's probably about time that we tackle it. But one of the important things you said was we've been doing this forever anyway, it's just the root server operators have not disclosed to the rest of the world what our metrics are for when we should report something or not. That after major events and things, we are always publishing notes saying this is how the entire RSS handled this particular—I remember when the last DDoS event was. It was quite a while ago that took out a good percentage of the root server system.

It's nothing recent. To me, though, one of the other questions besides what you put on the Board is, how do you define what things that we expect each root server operator to do, or do we leave them to come up with their own internal definitions for what reporting requirements are? Versus how do we deal with all of the root server operators in a joint coordinated effort for something that is reaching beyond a single operator? That's where I think if RSSAC wanted to get in the game, that's more likely where we should be. But I suppose RSSAC as a policy or ICANN as a policy could specify these are the minimum reporting rules that everybody knows is subscribed to in order to play the game.

KEN RENARD: Yeah, that makes perfect sense to me. That's one of the sort of scope of incidents or reporting I'll mention later on. Again, do we want to define this and what are some initial thoughts on these definitions? Exactly those types of questions. So thank you. Anybody else in the Zoom have any initial thoughts on what we should do? Starting a work party? Paul?

PAUL HOFFMAN: Hi. So before we have much more discussion, I think one of the things that Wes just brought up is very important because it will greatly change the expected outputs. I mean, we don't really want to start a discussion without knowing what the expected

outputs. And so to me, as somebody looking at this new, I would say that the most important question here is, are the expectations only things that would be reported by the whole RSS as compared to an individual RSO?

So for example, T-Root has a DoS attack against it. That DoS attack is not seen at other root server operators, maybe only one. Is this something that is supposed to be reported just by T-Root, or is it going to be reported by the whole root server system? The reason this is important is many root server operators use their infrastructure for things other than root service. Some of them are other TLD servers. Some of them are doing other things. So an attack on an individual or a small number, which doesn't have a material effect on the users of the root server system, is that something that this will affect or not? So with without knowing that, then we are saying everyone should be disclosing sort of in a general way. With saying that, we have to do some coordination. So it would be good if we knew that before we jumped in. Thank you.

KEN RENARD: Thanks. I think one of your phrases there, "material effect on the RSS," that seems to be a good starting point for where that line is drawn. If you have a single instance that crashes, that doesn't seem like a reportable event. So we can go on maybe to the next two slides. Sorry, Brad? Thank you.

BRAD VERD: I'm trying to think of the right words here. I feel that certainly somebody, maybe we need to define what the bar is as to what the reporting level is. I certainly don't believe Paul, in your example, that that reaches any level of reporting. There are attacks all the time, different levels, little to no impact, and I don't think that is something that would be worthwhile. I believe—and I'm putting words in just what I've heard from different people—is that we need to report on the signal and not the noise. And I feel a lot of what you just described is noise.

KEN RENARD: Thanks, Brad. Russ?

RUSS MUNDY: Thanks, Ken. One of the things that I think needs to be given careful consideration as part of this process, too, is what is the sort of end object of the reporting itself? As Wes pointed out earlier and those of us that have followed along this for a long time know there's always been a lot of discussion at the Root Ops level mostly as things happen that affect the root server ops. But if the objective is to try to have a set of criteria and information and thresholds that are reported just amongst the root ops or amongst even the RSSAC or maybe the ICANN, but things that might be considered just internal. It would be

appropriate to try to develop some structure and criteria to use versus things that would be reported to the public and made public and whatever the in-between might be, whether it was some government entity demanding information about something that happened or how broad the extent of the end report is intended to be. Maybe I've missed it but I haven't heard that really discussed. Thanks.

KEN RENARD: Thanks, Wes.

WES HARDAKER: That was Russ. I'm Wes, though. Thanks, comma. Right. There's a comma in there. You don't know the number of times that Russ and Wes have always had their names reversed. Russ can speak to that as well.

Let me make a solid proposal. The proposal is that RSSAC should only be defining what we're going to do based on the root server system, that in order to preserve the independence that we have argued for in both 037 and 058 and the GWG listen and there are documents too that it's up to individual organizations to do appropriate reporting for themselves. Now, the question is how do we define what affects the system? We know for a fact that the root has never gone down. We know for a fact that even when a significant portion of it has been

affected, the DNS system as a whole has been just fine. Maybe slightly delayed for delays that nobody will notice. But I would propose we just stick to what are we going to define for the whole service and we just drop the notion that we're going to work on individual operations that's up to each organization.

KEN RENARD: Thanks, comma, Wes. Liman?

LARS-JOHAN LIMAN: Thank you. I agree with Wes. My thinking about RSSAC's role here is to design the future, not to be reactive to what's going on out there in the realities of—it may sound a bit strange but I'll try to explain a bit better. To me, the operation is something that's handled by the roots server operators, and that includes dealing with various types of cyber attacks. RSSAC's role, in my view, is to look at the whole system and to design the system so that it provides a good service for the general Internet and so that it doesn't contain properties that make it difficult to defend in cases of cyber attacks. So I would argue that if the systems need to be changed based on what we see in attacks, it's up to the root server operators to bring to RSSAC a problem and suggest a change, and RSSAC can discuss this and take a wider picture than just the root server operation and build a better understanding of the situation and possibly, hopefully, do what the root server operators suggest or something along those lines

to improve the system as a whole. But I don't think that RSSAC should be burdened with operational things like reporting statistics, attacks, and so on. And I also think that Wes is right in that we need to maintain our independence from each other. Thank you.

KEN RENARD: Robert?

ROBERT STORY: My question is, do we think that the GWG is going to accept that we're just considering things as the RSS as a whole. I kind of get the feeling that they want to be able to hold individual root servers accountable. And maybe somebody that's more involved in that group can speak to that. But based on my impression that they are interested in individual responsibility, I think that it would be a good exercise for us to prepare for that because we have less representation there, and that if we can go through the exercise in RSSAC and get all the roots involved and start to figure out what we can individually as organizations do or want to do or able to do so that our representatives or the GWG can represent what we feel is realistic.

KEN RENARD: Okay. Just a clarification that the GWG has not asked this of us, this is something originally coming from the success criteria. It

would be something that the eventual governance structure would do—or maybe not imposed but would facilitate. I think Liman was next.

LARS-JOHAN LIMAN: I think I disagree to some degree with you, Robert, because I see the current RSSAC—current RSSAC, I'll stress that—not be part of the future model. There might be something RSSAC like but that will have a different role and will have different strings attached to it from different sources. Currently, RSSAC is an ICANN body that advises the Board, and I don't think we should engage in activities that the a future body committee, what have you, might do in the future because we get the strings attached immediately to ICANN. We may not want that. So to plan ahead, to look at what would the future bodies and relationships look like, that's a good thing, but we actually have the GWG to do that. I am happy to discuss it here in RSSAC as well but I would argue that the GWG is the main focus point for those discussions. Thanks.

KEN RENARD: Certainly. The governance structure should have this. We, as RSSAC, have already stated that. This would be RSSAC expressing a recommendation or advice to the GWG to then make a decision. Jeff?

JEFF OSBORN: Thanks, Ken. I'm not sure where this fits in. But I feel like it needs to be said. I really agree with Wes's comment that when there is an individual incident or an individual RSO, then the need for that to be broadcast widely isn't necessarily there. I'm just wondering whether there isn't a threshold event where if all of us are under an attack of some level and we hadn't thought through what do we do when we're all sort of struggling to keep the thing going, is there a threshold we could agree on in advance? I think people on the outside would feel better if we said, "You're not hearing from us because minor nothings are going on," yet another DDoS attack, yet another whatever. But letting people know that they're not going to just hear silence when an attack of level, whatever is deemed threatening enough, seems like something if we could figure out in advance, I think that would help the audience. And if I put that poorly, I apologize. I know what I mean. I'm just not thinking straight.

LARS-JOHAN LIMAN: Clarifying question please: hear from whom? You say there's a threshold. There's a threshold, and I'm fine with that. But what happens when we hit the threshold? Who talks to whom?

JEFF OSBORN: Good question. But somebody outside the 12 organizations it seems should need to know if we are at the point where we're answering one query per minute.

LARS-JOHAN LIMAN: So in theory, that would be the governance structure could make a statement, not an individual RSO. I go back to Paul's phrase, "Material effect on the RSS as the threshold." Of course, that means nothing with respect to measured bits, bytes, or anything. Certainly subjective but we can haggle the numbers later.

KEN RENARD: Yeah. Robert?

ROBERT STORY: Thanks very much. Just listening to the conversation play out, it seems to me that there are two different types of thresholds and two different types of reporting that are under discussion. The first would be the question, at what point do RSOs talk to one another? In other words, at what point does an RSO disclose to the RSO community? However that's going to be defined about a difficulty that they've encountered.

The second question would then become at what point does the RSO community, through whatever mechanism ends up being

created, disclose what has happened to other stakeholders in the process? I suspect you'll probably end up in a situation where you set those at two different thresholds where there's more internal reporting and information sharing than otherwise, but perhaps not. And then I think you've already mentioned in this last exchange, this other point about the actual mechanism of who was talking to who, because I just used the phrase the RSO community rather loosely. I mean, clearly GWG is in the process of discussing what body will exist that can receive, and then possibly onward transmit those reports. At the moment I'm guessing from what I've heard that there's a rather informal mechanism in place.

One preliminary thing you may wish to consider is do you have anywhere written down a taxonomy or a system for assessing when do RSOs talk to each other about what they're facing? Because clearly, it happens informally. And you might want to ask yourself the question, is there value in triggering a process where RSSAC fashions or begins to think about a recommendation about when is it appropriate for RSOs to talk to one another about an ongoing incident?

KEN RENARD: Thanks. Yes, RSOs do have individual thresholds of when they talk to each other and they talk a lot. I would almost think, hope that the consensus would be to have that not in scope of this.

This would be really more just the RSS. If something's happening day to day, RSOs can talk to each other at any point. So At least in my mind, I'm thinking this is more of the bigger picture and what's happening to the RSS. How does the root server system maintain transparency and enough to maintain the trust by the ICANN or Internet community in the root server system? I go to Russ.

RUSS MUNDY: Thanks, Ken. There's been a fair bit of discussion about the term that I think Paul used first in the material effect. One of the things that I have to point out because although I've not had that much direct involvement in the last few years in disclosure and things like that, I spent a lot of years in operational things where I ended up being right in the middle of it. I would say that if we're going to as a group or as a body, try to come up with something that is committed to paper and published, that we should try very hard to base it on our existing published performance-related criteria, RSSAC047, 002, as the basis and tie things to that, leaving room to say, "Oh well, there may be other things that end up that we choose to report that seemed to be really important." But I would be very opposed to trying to define some kind of subjective mechanism as the main criteria that's used because I can tell you, when you get in the middle of these things, it's almost impossible to know when you fell over that subjective criteria. So I don't know. Others may not agree,

but I think it's best to start with something that's tied to your existing number commitment for performance. Thanks.

KEN RENARD: Yeah. What do we as RSOs or as RSSAC or the governance structure think is important enough versus what does the community want from us as far as disclosure or reporting that's a good balance. Wes?

WES HARDAKER: Thanks, @Ken. I forgot to lower my hand. You said what I wanted to say earlier, which is the internal communications between the RSOs is really out of scope for RSSAC. But I agree. In fact, we have an internal policy of when we start reporting stuff to the rest of the RSOs based on our internal measurement metrics and things like that, and I everybody has their own. I was repeating what you said. I never took my hand down, sorry.

KEN RENARD: Yeah, most things start with—oh, that's weird.

WES HARDAKER: Is anybody else seeing this weird thing?

KEN RENARD: Liman?

LARS-JOHAN LIMAN: Thanks. I would like to respond to Robert here who asked, "Is there any value in specifying numbers and thresholds?" I would argue yes and no, where the no part is we don't need that for our operational procedures right now because it works. We know what the thresholds are, we know when we need to talk, and that actually works very well. But it may have a value from a transparency angle where you look at this as an ISO certification that if we specify numbers, we give people outside the root server operators' community a chance to know what to expect from us. So they will expect that if things happen above these thresholds then we will be notified, and that can be a good thing that transparency in that. Thanks.

KEN RENARD: Paul?

PAUL HOFFMAN: I'm going to respond to something that Liman said right at the end but a few people have said as well, which I think is another consideration that will need to be made, which is do the RSOs tell the community while a cyber attack is going on, or do they report afterwards about it went on, here's how we dealt with it, here's how it ended and such? Those are two very, very different requirements.

KEN RENARD: Yeah, good point. That's something that was brought up, I believe, in NIS2 context where thou shalt report within X number of hours, 24 hours. I think most of us, the hands-on keyboard folks are going to say, "I'm not going to be filling up forms here. I'm going to be trying to fix my systems." So what I see this is more defining what those might be, what the thresholds of, "Hey look, we're not going to tell you we're taking enemy fire right now. We're going to set some expectations. We'll tell you when we can, probably it won't be at a certain number of hours, but might be at least some summary or wrap up."

This has definitely been some good discussion, exactly what I was trying to invoke here. If we have no other questions in the chat or in the Zoom, we can go to a couple more slides ahead. Slide 9, discussion guidelines here. I think we're actually going to pass this here. We're definitely not going to resolve this topic or create a revised document today. So what we can do today is have enough discussion to inform or maybe make a group to build a statement of work.

These are some of the things that I had talked about. Some of them have already been discussed. Obviously, we need to retain RSO autonomy and independence. I'll go to Wes.

WES HARDAKER: I was actually backtracking a little bit because we actually never talked about—I think an important question you brought up is how do we deal with a subject? And you said the obvious choice would be the RSSAC Caucus. I 1000% agree with that this is the exact type of thing where we want a wide variety of input to think about and discuss and come up with the document. So I'd propose that compared to the other bullets you had, one of which, of course, was do nothing.

KEN RENARD: Yeah, the cookies are gone out there. So yeah. That's hopefully the answer we'll have by the end of this session, if that's the consensus even adding in the option of bringing in SSAC specifically. Is there expertise that SSAC would have that would be helpful that we don't have here? There's certainly SSAC folks on the caucus. As we looked through some of these things here, I broke them down. In the statement that's in RSSAC058, it talks about cyber incident oversight, it talks about disclosure as two separate things, and then information sharing with the governance structure and among RSOs. I think we've decided among RSOs, that's pretty much out of scope. That's something that that we do within Root Ops. So looking specifically at a governance structure that would do cyber incident oversight, what should that look like? I talked a little bit before about the boundaries of what's in scope and what's out of scope. So if an RSO's organization HR system gets broken in two, that's out of

scope, that's the obvious thing, only things that would potentially have a material effect. So these are things that are probably pretty obvious to us but maybe should be written down.

Again, an incident oversight, specifically, the purposes, we can share lessons learned. We do that already. Things that could affect potential architecture changes and then evaluation of an RSO individually to provide a service. These are some of the reasons that I thought up for incident oversight, and any thoughts specifically on our governance structure, whatever it might be, just overseeing security incidents. I'll go to Jeff and then Paul.

JEFF OSBORN: Thanks, Ken. In looking at the card on the screen, the thing I'm really curious about is the degree to which we do this internally has a long history and there's really just a matter of fine tuning to it. I'm wondering if you view this from the point of view of somebody who is outside of this organization and is interested in the idea that this group is becoming more responsible to something. So let's say I am one of the multistakeholder community and I am curious what goes on within this organization, what controls are there, how are incidents reported, and most of the things I'm looking at are sort of "none of your business, we're doing this internally, we've got it

handled, thank you," which might be the right thing. But if we stick with that's the message we're sending, we're going to recognize at some point, we're going to have to send the message of, "We're pretty good self-regulating. Thank you very much. Have a nice day." You know what I'm saying? At what point does this start to sound a little parochial?

KEN RENARD: I get what you're saying. That's what I would like to almost have not written down or just at least refer to and maybe have us define it rather than it defined for us. Again, if we saw a two megabit per second spike today, I don't want to have to fill out a form. We can start to define some of those thresholds or what it means to disclose.

JEFF OSBORN: Right. To sort of fill in the word you didn't exactly say, but you kind of said, "Do you believe that if we don't do this, somebody else is going to request it?" And that would be harder.

KEN RENARD: I don't want to have that pessimistic of an attitude. But if I had to look at this from the perspective of what should we as this critical infrastructure, what do we owe to the Internet community as far as transparency? Paul?

PAUL HOFFMAN: Thank you for that immediate previous discussion, because one of the things that many of us when you talk about this community, I think, Jeff, you are probably talking about the root server operators. There's the wider community of the caucus or people who follow and such like that. And many of us have heard informally the RSOs talk with each other about incidents, they have a back channel and such like that. I believe simply stating that in an RSSAC document that this not only exists but it has existed for decades would go a long way towards relieving some of the concern of, "Oh, do they need our help in being able to coordinate with each other?" The answer is absolutely not. In fact, they're doing it better than the people who think they can help.

The reason I raised my hand, though, is, Ken, I'm very, very concerned about the third bullet under purpose of incident oversight. Evaluation of an RSO's ability to provide a service I think it should be absolutely out of scope, both for the independence issue and also because I don't believe that we as a community understand an individual RSO's operations well enough to help them to evaluate how they are providing the service. I would hope that if we take this on, and I think we should, there's nothing in there about helping individual RSOs, expecting reports from individual RSOs, maybe this RSO is

considered to be needing more guidance than others. I want none of that in there at all.

KEN RENARD: In this work party, I certainly agree. I think in the sense that the governance structure will oversee the performance of RSOs, their ability to provide the services is affected by cyber incidents. But I can definitely see your point and allow that to be defined more in the designation and removal versus the specific document or work party. These bullets are here really to spark discussion as thoughts. I'll say something controversial just to get the discussion going. I'll go to Kaveh and then to Liman.

KAVEH RANJBAR: Okay. Thank you very much. I get a point that it immediately gets basically to a value judgment or is even a cyber incident. So I would like us to maybe take a step back and rethink this because we published the RSSAC002, and that shows a number of queries we have responded, what's bandwidth, things like that, we provide overview on few metrics. And of course, we can amend that if needed. If you look at the system and keep in mind independence and all of that, we are 12 independent organizations. Together, we provide one service, not single one of us is responsible for that. So a collective of us provides that service. In doing so, anybody with the information that we provide is actually eligible to look at how we perform as a

system, I think that's an important part, and judge us and then provide feedback. Actually, I think it's much more healthy if someone from outside based only on the data must decide that, "Okay, they had suffered an attack" or something like that. So internally, that will be up to me if I have right incentives to provide my service, protect against hacks, and provide answer to as many queries as possible. I think that's the healthiest model, that I follow my own incentives, and then do the best that I can. Then someone else can look at the collective and say, "Hey, the collective is not doing well or doing good or this thing is getting more or getting less." I think that would be the most sustainable approach because everybody is doing what's they're being paid for, basically.

KEN RENARD: Right. So that autonomy and independence. In your mechanisms, what specific configuration you have on your DoS protection, we need to preserve that independence. Liman?

LARS-JOHAN LIMAN: Thank you. I have three comments. The first one is that we probably need to realize that things are going to change. In a couple of years' time, we're not going to operate in the same way that we've been doing so far because we are creating this new structure, on that I predict will impact the ways that we work and cooperate in some ways. I would prefer it to be in good

ways. We need to think outside the box a bit and the box being how we do things now. In that respect, I think that there may be channels where we will be expected to talk to each other under certain circumstances and that we possibly need to be transparent about which levels these are and how we communicate in these cases. But that's something to be decided in the future when we start to look at the performance monitoring things and we should definitely be part of defining that ourselves.

Second thing is that we do a lot of operational talk between ourselves already. But if we look at the future where there may be new root server operators designated, I think it should be a requirement to participate in that type of collaboration in the checklist for new RSOs, so that we don't have a single new RSO who doesn't participate in this working social system that keeps things running.

The third thing is that maybe handling incidents should also be either part of the SLA or part of checklist when a new root server operator is designated so that there is an expectation for how incidents are handled and there is an expectation to participate in that work and collaborate with the others. That it's not just about performance numbers, it's also about social interaction in the requirements. Thanks.

KEN RENARD: That's one of the RSO principles about participating in the community. I whole heartedly agree. Russ?

RUSS MUNDY: Thanks, Ken. Kaveh brings up an interesting way to express this. In some ways, I think, if I understood him correctly, it's quite similar to what I think mentioned earlier and that if something of this nature gets written down, it should definitely be tied against the existing or evolving performance requirements that are already specified. But if it's done in such a way that—and I don't know our information that we have yet from our monitoring, especially related to our RSSAC047 because I haven't been watching how much output and the utility of the output and so forth. But if it's mature enough, in fact, it would be practical to essentially say that this is a result of RSSAC047, the monitoring and the data collection. And if it passes a certain type of threshold that gets identified, then it will be reported as a cyber incident, as opposed to saying, "We're going to have the cyber incident reporting, blah, blah, blah," and then try to define what the material effect on the DNS means and the many other subjective criteria. Because I think, Ken, your suggestion to write this down first ourselves is probably a good one, because if we don't, somebody else will come in and do it for us or tell us we got to do it in an unreasonable timeframe. But do it in a way that tightly relates to things we've already defined. Thanks.

KEN RENARD: Thanks, Russ. All right. Maybe go to the next slide. I broke out to the oversight by the governance structure versus disclosure obligations here. Just a few thoughts on that. Setting expectations for external organizations. Again, if the general public knows that if the DoS attack is this tall, it will be reported. And if the reporting system is quiet, they know that the current DoS attacks are less than that. That can be certainly good for transparency and improve trust of their reserve system. But some additional thoughts on disclosure. Again, we shouldn't be filling out forms for every spike in traffic and certainly don't want to leak any information that might be useful to an adversary. And again, be respectful of individual organizational policies. We each have our own publication, processes, things like that. Again, going back to Paul's material back in the RSS, I think that's a good threshold. Go to the next slide.

This is, again, breaking down that original piece in RSSAC058, the individual pieces here, the security and vulnerability sharing among RSOs, best suited there for Root Ops. Within the governance structure itself, can we set those thresholds of what should be reported? And number of RSOs saw this level of DoS attack. Again, how do we find that material effect on the RSS? What should the governance structure then do as far as making a statement? All right, so those were some just thoughts. Paul, you had your hand up.

PAUL HOFFMAN: I'm going to channel Brett Carr here because I don't think that people may be watching the chat. Brett brought something up earlier that I think is fairly important, which is this is not all about DoS attacks. There could be other attacks such as software attacks and such like that that could have a material effect. I think it's okay to have DoS be one of the major parts of the discussion, but it should not be the only part.

KEN RENARD: Absolutely good point. Brett?

BRETT CARR: I'll just back up what Paul said because I brought it up originally. There are lots and lots of things that could be cybersecurity incidents and we need to not focus on just DDoS and thresholds because those thresholds and those incidents could be very, very different to DDoS.

KEN RENARD: Absolutely. Anything that would have that material effect. Either if it was changing records or somehow changing the accuracy of responses, things like that.

BRETT CARR: And changing the risk level.

KEN RENARD: Exactly.

BRETT CARR: Obviously defining material effect is quite important as well.

KEN RENARD: Okay. Well, this has been a very good discussion. I really appreciate the input. At this point, we can go back to the original question of what does RSSAC want to do with this? Do we want to spin up a Caucus Work Party? Is that the right place to do this? I think a lot of our discussion captured here can either go into a statement of work or into the introductory material for the potential work party. Liman?

LARS-JOHAN LIMAN: Thank you. I think we should first answer the question, what's the expected output? When we know what to do, then we should address how to do it.

KEN RENARD: I guess the expected output would be a document that expresses the RSSAC's advice to the GWG on what we think a cyber incident oversight and disclosure means within this

context to encompass any or all or even more of what we've talked about here, and setting expectations what we think it should look like and then passing that on as just advice to the GWG. Kaveh?

KAVEH RANJBAR: So if I can frame it based on what I tried to mention before and Russ's comment, I think actually there is a statement of work in maybe defining incidents for the root server system. So what would constitute an incident for root server system based on any metric? If that percent of resolver received incorrect answer or invalidated signature or whatever the delay or lack of access, things like that. So if we can define what constitutes from our point of view, an incident for the root server system, I think that would be the best entry point, and then GWG can do whatever they want underneath of that. So that's what I suggest. If you want to define basically work, I think this will be the most useful.

KEN RENARD: Go ahead, Liman.

LARS-JOHAN LIMAN: Thank you. I think it's a good idea to write a document as input to the GWG. I would suggest—I haven't really wrapped my head around that yet, but I think what we want to do is to give them advice that a cyber incident should be a thing that's baked into

the final solution and to ask them or give them advice on which information channels to build into that system. So that information can be channeled from the root server operators to the general public in a good and predictable way. Also to define in the structure of point where such information can be collected and assessed for larger incidents. And possibly even to make sure that there is a corner where collaboration between the root server operators in these incidents can happen.

So that's I think what I would like to see as output for this. Not try to define what's an incident yet because that's further down the line, I think. But to back out and see, we need these channels, we need these methods of collaboration, we need these ways to publicize data or information and make sure that there's room for that in the future model. Thanks.

KEN RENARD: Wes?

WES HARDAKER: I think I agree with Kaveh that narrowing it down to what's the starting set of material is a good way to go. And it seems like the starting point of what are the things that the community will likely care about is a good one. I'll take note that I think we are misrepresenting to some extent sometimes what the GWG is doing. We are in the state of a huge period of flux, right? What's

going to come out of the GWG is not a complete governance system, it's a bootstrap. It's how to get started. I'll let you talk more about this on Thursday, in fact.

I don't know that if we gave this input to the GWG, they're going to have the time to put it into their model. And this will be something that some part of the GWG will do in the future for creating new policy. So it may be an input but I don't think that … Let me put it another way. I don't think they should actually take that in and start working on it immediately. It's some sort of input to a future body which is yet undisclosed.

KEN RENARD: Okay. Valid point. And if this work should be done at some point, if we can get a head start on it and help, that's also a good thing. We are running to the top of the hour. What I hear is a consensus. I want to validate that. Does anybody think that we should not start up a RSSAC Caucus Work Party on this topic? Go ahead.

UNIDENTIFIED MALE: No, I'm not going to disagree with that. I was just going to say it sounds like there's an action item for staff here to start a statement of work and probably work with you on that.

KEN RENARD: I would open up the floor to anyone else that wants to join to write a statement of work.

ANDREW MCCONACHIE: Sure. I mean, it will eventually be shared with the caucus, but I'm willing to work with anyone. Brad, go ahead.

BRAD VERD: I don't know. I can't put my finger on it, but maybe a work party is too soon. I feel like maybe we need to have a couple more discussions on this just as a group to work things out before we try to sit down and capture our thoughts and the statement of work. I don't think we're in a position where our thoughts are clear enough to do that yet.

WES HARDAKER: I'm in the queue so I'll just go ahead and jump. I think you're right, Brad, that we need … I think everybody wants to do this but we have to define the scope carefully, and that means defining that what goes into the statement of work carefully. I don't think we're going to turn the statement of work around in a week. You're absolutely right.

BRAD VERD: Which just means more talk. That's all.

KEN RENARD: Should we take the generation of the statement of work to the RSSAC Caucus mail list or the RSSAC list? Liman?

LARS-JOHAN LIMAN: I'm very much on Brad's side. You hit the nail. We should not take anything to the RSSAC Caucus list just yet. Maybe when we have a strawman or something that when we have collected our thoughts more, but not at this point, in my opinion. Thanks.

KEN RENARD: Okay. Andrew, you and I could do a very rough draft and then have a statement of work, and then circulate it among the RSSAC just to scope this. Russ?

RUSS MUNDY: Thanks, Ken. That was going to be my suggested next step, is that maybe a very rough cut at it at just on the RSSAC list. And then if the RSSAC feels we're far enough along—because I'm in agreement with Brad and Wes that we aren't far enough along yet about what we actually do want to accomplish. So let's talk some more on the RSSAC list.

ANDREW MCCONACHIE: Sounds good.

KEN RENARD: Okay. At the very high level—Ozan, if you can go back to slide four. Going back to that excerpt from RSSAC058, most of the governance structure must include a provision for cyber incident oversight and disclosure, etc. At a very high level, we want to define what this means. It sounds like a statement of work needs to get more precise before spinning up a work party. With that, we're a few minutes over. I thank you very much for the discussion. This was exactly what I was hoping to spark. We'll send out something fairly soon to discuss further statement of work. With that, is there anything else you wanted to say, Ozan?

OZAN SAHIN: No. Thank you, Ken. Thanks, everyone, for joining today. We can stop the recording.

[END OF TRANSCRIPTION]