ICANN75 | AGM – RSSAC Work Session
Saturday, September 17, 2022 – 13:15 to 14:30 KUL

OZAN SAHIN: Hello and welcome to Root Server System Advisory Committee Working Group Work Session. My name is Ozan Sahin and I am the remote participation manager for this session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior. Please note that this session is intended for discussion among the Root Server System Advisory Committee members. Other participants will be silent observers.

If you would like to speak during this session, please raise your hand in Zoom. When called upon, virtual participants will unmute in Zoom. Onsite participants will use a 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 this session in the Zoom toolbar. With that, I will hand the floor over to Ken Renard.

KEN RENARD: Thank you, Ozan. This session is split between a discussion of the cyber incident reporting as well as discussing RSSAC000v7. So roughly half the time will go between the two topics and adjust as necessary.

So if we go to the next slide. It's probably slide four. While we're getting there, the cybersecurity oversight and disclosure. This was

based off of RSSAC058, the success criteria, which states in there that, "The RSS governance structure must include a provision for cyber incident oversight and disclosure obligations and codify security threat and vulnerability information sharing amongst RSOs and the RSS governance structure body."

So it's important to note that this is something that has been ongoing. RSOs have been doing this type of work, this information sharing amongst as well as to the community, since the inception of the Root Server System. The motivations here for doing this are for providing transparency of the Root Server System and its operations, to continue to earn trust in the Root Server System from the community, and also to share lessons learned. So this bridges the gap between the ICANN world and root ops in some senses. But we're trying to figure out what can and should the RSS governance body do to facilitate reporting.

So this is a proposed work party. It's not something that's been kicked off. We do not have a statement of work yet. That's, in fact, one of the things that's holding this up. We had a draft statement of work and there's some ambiguities in there. There hasn't been much in the way of comments or discussion of the content of that statement of work yet.

So the idea of today's session is really to throw out some ideas, start some discussion, and see what we come up with that might further help us define what that statement of work would be, whether this should continue, should go forward. This is not something that we need to pursue. A work party would actually be making a

recommendation to the GWG or the RSS governance structure on the topic. So we don't get to say what happens. We get to just make recommendations.

So the topic of cyber incident reporting is very complex. The issues of independence and autonomy of the RSOs is very important to us. So each organization is going to have their own security policies. So doing something broad and applicable to the entire RSS is going to be complicated there.

And we need to balance things like transparency and security. Transparency, the more information exchanged is better. Security, sometimes the less information is better. So how do we balance those things?

In the statement of work, there was a good attempt at trying to define what these disclosures should mean. One of the suggestions was to disclose things that would have a material effect on the Root Server System. Definition of material effect is very open to interpretation. What does that mean? And how can we define it? If we're going to put together a statement of work, we need to at least have some sort of bounds on what this might mean.

So if you go to the next slide, what we wanted to do today is just throw out some ideas, and have a brainstorming session, and then collect the results of that to try and further the statement of work. So the topics to banter around here today are to explore the perspective of the consumers of this disclosure. What would a consumer of this disclosure information want? What position would they be in that

would make them want this information and what pieces of information would be useful? And then, what roles could the governance structure take? So these are all open-ended things up for discussion.

If we go to the next slide, my thoughts on the consumers of the incident disclosure. Three levels of disclosure there—public, among other RSOs, and with the RSS governance structure. So different domains there. We're mostly RSOs here. We're mostly in that second open bullet there and probably going to sit, eventually, close to that third open bullet.

But what—Internet users or a proxy for what is good for Internet users as far as cyber incident disclosure. Regulatory bodies, governments, what types of perspective would they have? What type of useful information would they get out of cyber incident disclosures on the Root Server System? How do we engage these communities? Should we and get their perspective? If anybody has any thoughts what these bodies are, what these governments are, what they would want … Sorry. I've not been monitoring that. Yes.

JEFF OSBORN: One of the first things that comes to me here is who is requiring that we provide this information and who is determining what the audience for the information is? Because the most obvious other people who would want this information, and for whom it's useful, and we can trust it's used well is the second two bullets—the second two empty ones.

To the degree that this is public information, it just seems like if we're the ones who have to decide that, I'm not sure mistakes weren't already made. We released—I'm not even sure if I'm allowed to say it but a fairly large amount of CVEs are in process with BIND right now. And that tends to be something awfully secret involving the public. And we very rarely, gratuitously, let the public know things without having gone through quite a checklist.

So when I see this, I'm just curious where is the demand or the need that cyber incident disclosure immediately is a public thing on par with everybody else there.

WES HARDAKER: I was going to say something different but now I have to answer Jeff. Speaking of rules of adults, I think modern framework, the whole reason that you're releasing CVEs for BIND on a regular basis is because it's the right professional thing to do these days. And I think, certainly, running a critical infrastructure service for the Internet, the right thing to do is disclose when things happen.

You bring up an interesting point, though, which is about timing—the timing of what is acceptable for how long to wait before disclosing that something happened. Do you have to complete remediation first? I think all of that leaves something to be desired, of course. But in the world of new adults, I think releasing to the public is probably something that we should do.

The other list of who this might be of interest to, it occurred to me, actually doesn't exist yet but it will in the future, which is the outcome

of the GWG and the auditing department. So at some point … I don't remember what the circle is on the map but that circle will certainly want to track us for how well we're doing.

KEN RENARD: Thanks, Wes. So in my mind, the RSS GS is whatever has that body of auditing or whatever that would "oversee" the RSOs. Yeah. So why? Why would we disclose to the public? For transparency, to maintain the trust in the Root Server System. Hopefully this, for many years, will just be a no-op, like, "Nothing interesting happened this month."

But if I were an average Internet user, knowing that here is the threshold, knowing that nothing happened, no material effect on the RSS has caused a problem recently, that's good. Let's say, if there was a threshold—if the attacker was at least this tall or the attack was at least that tall—knowing that nothing smaller than that had happened recently would help put me at ease. Jeff?

JEFF OSBORN: Should I, then, not be concerned that we're trying to address the incidence of, "The EU noted a blip yesterday. At blip plus two hours, we're looking for a report that explains what it was?"

KEN RENARD: I personally don't think so. I think that if this goes through a work party, they're the ones that would probably determine what size blip. But going back to a material effect on the RSS. So there's a lot of fine-

grained stuff out there. An HR system of yours got broken into. That seems ridiculously out-of-scope. But go to Terry.

TERRY MANDERSON: We have been publishing formulaic incident reports for ICANN IMRS for over the last year and a half, roughly. We have not yet had any requests about any blip or any kind of events from any governments. We have certainly heard of governments wanting more information. And you might was to consider this, perhaps, from my personal opinion—not ICANN's position but my personal position. This is a preemptive retaliatory strike to overregulation by other entities out there.

So if we publish it and we act as, as Wes described, an adult in the industry, then it's going to build the trust. So I don't actually see this as a bad thing. What I am doing, though, is treating all circle points on that list as exactly the same. I don't differentiate. I don't give any special more information to the public than I do the governments. So there is no special source there. Hope that helps from my perspective. Thank you.

KEN RENARD: So there's no distinction between anything in that public?

TERRY MANDERSON: None whatsoever.

KEN RENARD: Okay. But maybe, to other RSOs, it would be different.

TERRY MANDERSON: That's a different conversation. They can still see the reports. Other RSOs can still see those reports. If two RSOs engage in a confidential discussion, that's a completely different thing.

KEN RENARD: Liman?

LARS LIMAN: I have two things that I think we should keep in mind when we look at this and develop it further. One is how do we perceive that the various bodies here are going to use the information? Is it going to be basis for decisions made by whom for what purpose?

Let me see if I can remember my second one. Yes. If we were to receive these blips and questions regarding blips, I hope that, in the future, the outcome of the governance discussions structure to which these queries can be directed, which can shield the root server operators from the really stupid ones that can be dealt with really quickly. But that's maybe wishful thinking on my part. But that could be something that we could design into the system if we want to. Thank you.

KEN RENARD: Wes, go ahead.

WES HARDAKER: If you look at the corporate world, the corporate world does not have a standard reporting level for everything. I don't think that we need to go down that role either. I think what's more important is that RSOs should document what their level is. And by documenting it, then somebody can complain if maybe, "$3 trillion of damage is a little too high of a bar. Do you mind lowering that a little bit?"

Internally to USC/ISI's root, we actually have some levels of which. We haven't documented them publicly. But internally, we're like, "Okay. If we're down on DNSMON for 20% loss for 20 minutes, then we tell at least somebody." Terry, I appreciate your root's particular publication of your regular reports because we're probably going to use a lot of that as a template.

But the important thing is, I think … Terry's root's significantly different than mine. The important thing is that we document where those levels are and what we surpassed because that allows feedback from the community of is that sufficient or not.

KEN RENARD: Right. The question will come. Should we be talking about this for the RSS or individual RSOs? I'm thinking more along the RSS, whereas we've already stated that the failure of an RSO or RSI is not a problem. So if one root server went completely down, is that a reportable event? Certainly, as an RSO, you've got some explaining to do. But that might be a different community, then. That might be discussing among the RSOs versus a full public disclosure if a significant event happened to the entire RSS. So, Jeff then Terry.

JEFF OSBORN: I think a good question is who is being compelled for this information? I'm a huge fan of what Terry is doing. And I love the idea. It could be a baseline. If everybody just wanted to pile on and go, "Let's all provide some baseline something." But you make the good point of shouldn't that be at an RSS level rather than by RSO? So if nothing else, it gives us an easy way to say, "We're certainly willing to provide some baseline across the board."

I'll note, though, that RSSAC058 specifically only talked about reporting, l believe, between RSOs rather than to some public. So that might be out of the scope of how far we've been directed so far. Again, that's why I'm saying who is compelling this information? Why are we feeling the need to put this out to whom?

KEN RENARD: So the statement in 058 … I don't know—Brad, were you the one that put that in there?—if it was meant specifically between the RSOs and the governance structure? Certainly, we could think of this as what do we, as the RSS, owe to the Internet community as far as just disclosure, establishing trust, and being transparent?

BRAD VERD: I think that criteria was put in 058 as a response to the different regulatory or government structures out there doing something. Rather than having each of the RSOs respond to multi different regulations from different nation states, or communities, or anything

else, the intent or the goal was that the RSS governance system would be a single point of contact for it.

So, Jeff, when you ask who is compelling you, I think that's a hard question to answer. It's the cart before the horse, if that makes sense. You're trying to answer the question that hasn't been asked yet. And I think the question here is about disclosure and what are the RSOs willing to disclose.

I really don't think that the question here is who's going to consume it and what they're going to use it for. I think once it's disclosed, it's disclosed. People can use it for whatever they want or whatever it's going to be used for. I feel like once it's out there, it's out there. But we, as a group, going to what you said about Terry's disclosure that the IMRS does, is that it's out there.

JEFF OSBORN: I have erred. I was trying to make this simpler.

BRAD VERD: I'm all about simple.

JEFF OSBORN: Great. So if that sounded complex, then I'll withdraw it. My intention was I just wanted to make sure we knew why we were doing what we were doing and then we do it.

BRAD VERD: Yeah. I think, going specifically to your question, the success criteria in 058 specifically says, "Disclosure from the RSOs to the RSS governance system." You are correct. And then, obviously, the next step in there, somewhere there is going to be a regulatory body, or a state, or somebody who's going to want to see those disclosures. But they wouldn't be coming to the RSOs. In theory, they would be going to a governance system that has yet to be put together.

KEN RENARD: Terry?

TERRY MANDERSON: I think there's some nuance in what Brad's saying there. And I think that we're reaching a point where we need to understand from where and to whom we're actually going to be doing the reporting. I think that's the very next step we need to be looking at. And if it's the RSOs, it's RSOs collectively. So it's then completely about the RSS, not the RSOs collectively about one other RSO. I think that's somewhat nuanced.

And that changes even what I do in terms of my formulaic monthly disclosures. It will be something that one would conceive that a secretariat-like body would convene a discussion and then push out an appropriate disclosure if there was a significant event against the Root Server System. But I think we need to clarify that because the words there, what I just said, differ to the 058 wording. And I think because in the 058 wording, we're talking about RSOs disclosing to the

RSS governance body, not RSOs collectively disclosing to some omnibus public space.

BRAD VERD: Terry, I absolutely agree because 058 is success criteria for a governance system. That's why that tie is there. I think if something were to be done ahead of time then that would have to be figured out. And I don't think it's beneficial. I think it would be more beneficial … Let me rephrase. It would be more beneficial to report on the RSS rather than individual RSOs.

So while I applaud the IRMS for what you're doing and I think that's great, I think we, as a collective—the RSOs and RSSAC—have spent a lot of time and effort to inform the community that looking at the individual identities of the root … It's more important to look at the RSS as a whole. I shouldn't say as a whole. But the RSS, the availability of the service of the RSS, and not the availability of the individual letters. And we've shown that in the individual documents that we've done and other things we've stated. But I think disclosures would take a little bit of effort that the RSOs are currently not set up to do.

So I think you are correct, meaning there's a nuance there that would have to be managed if we wanted to do it ahead of time. If the work here is to do something to feed the RSS governance system, then that's a different discussion. It's not a different discussion. It's just we don't have to have the nuanced one of what we do in the meantime.

TERRY MANDERSON: Agreed. I'm just looking for what that next step is. I feel like we're at a juncture right now that we're asking questions about consumers. But we don't have a strong understanding of where we're going to position ourselves as a governance system. Are we approaching this topic too early in the frame?

BRAD VERD: I think so. I think that question is too early. I don't think we should be worried about the consumers right now. We should be more worried about the disclosure and what should be disclosed and what should be talked about there.

TERRY MANDERSON: Thank you. Exactly what I was thinking.

KEN RENARD: So my intent in talking about consumers—go to the next slide—is to find out what types of information. There's a lot of information, security-wise— "I updated this patch. I installed this piece of software—" that has no business going out to the public. But if there's a significant type of event that happens, what would those be? Where can we draw that line of what type of information, nothing specific that should be disclosed.

Then go to the next slide. There we go. These are just some ideas of what information might these different groups want. So thinking about it from the security perspective of some sort of loss of data integrity. What if distribution keys were compromised? Is that

something the public should know? Or anything that results in incorrect responses. If we had, maybe, severely old data or just giving out wrong signatures or something like that, that seems like something reportable.

Loss of availability, whether it's internal to the Root Server System or some external event that happened. Maybe a route attack. What about privacy violations? What type of reporting would be necessary for that? Which brings up the question what are privacy expectations of users of the root server system?

So it's just looking at it from a different perspective of these are some potential major effects or major things that could happen to the Root Server System. Are these worthy of disclosing to the public? Can we think of other doomsday scenarios where something bad happens? What would an Internet user or a regulatory body want to know? What kind of doomsday thing would happen? Brad?

BRAD VERD: I'm not sure it would be very beneficial to try to iterate through all the doomsday scenarios. I think it might be more beneficial to, I feel like, whoever these consumers are of, "What does business as usual look like?" Then, if business as usual all of a sudden changes, maybe there's a disclosure as to why versus trying to iterate through every exception, which I think will be hard and you'll never get them. That's all.

KEN RENARD: Right. Trying to figure every doomsday scenario is futile. But looking at it from the security principles of integrity and availability, things like that. Wes then Jeff.

WES HARDAKER: Thanks, Ken. I don't think that's the intent of this slide. Or at least my guess is that it wasn't. But I actually love that first example. And I love it because it is actually, I think, the best example we've run across that would affect the entire system.

So I don't think that we need to iterate through each one. But having some examples to discuss as to what percentage of the system was impacted. The first one's nearly 100% because we don't know how many wrong answers were potentially given or something like that. It'd be hard to trace down. Whereas the other ones, the [inaudible] RSI is that gets a percentage of the entire system and things like that so understanding the different levels does become important for how do we document where the line is between a single RSO being affected versus the system.

And we have, for the system, published things on rootservers.org for past large events. They happen very rarely. And we've done it occasionally, saying, "Hey. This wasn't just a single entity. There was a major DDOS attack that took out a significant portion of root. Didn't fall over. It never has. But you might want to be aware of it." So that's the incident reporting that I think that that first bullet becomes important to think about.

KEN RENARD: Jeff?

JEFF OSBORN: Every time I see a page like this, it concerns me. This is like googling your own health issues. If you have somebody who knows very little about how the system works, they could be concerned about a whole bunch of bad things happening. The actual end result is the root server system didn't even hiccup. But thousands of viruses were attacked by thousands of antibodies, and at the end of the day, everything's fine.

When you bring them out, I especially worry in the hands of people like governments, and bureaucrats, and the rest of them, being able to say, "There were 42 incidents of nodes going down for more than eight seconds," the end result of which we all know is nothing. But you worry that people who are unknowledgeable could just make a hash of this and make it really difficult to report anything meaningful because you get so lost in the details. Maybe this is unnecessary and we all know this. But it just scares me every time I see a page like this.

KEN RENARD: Right. And what I was thinking with his is just if … Loss of availability. There would be some threshold that the work party could come up with, x% of the entire RSS or whatever. If that threshold is crossed, we would notify somehow, some public statement maybe.

And then, from an Internet user, "I know nothing's been reported. Therefore, x% of the Root Server System must be available." That's that transparency, hopefully a good, warm, fuzzy feeling that things

are all right. But yeah. I would not expect that any of this would be, "0.39% of the RSS nodes are down right now." That just does not matter. That's just normal patching and updates.

Anybody else have thoughts on this? What about privacy? What types of privacy information should be shared? Is that even a concern? Brad?

BRAD VERD: I was going to go previous to privacy. But really quickly, on PPI, I don't think there's anything in the root that's PPI. So I have a hard time with this question as it's stated. I think Wes, going to your comment earlier, I'm trying to think of how to make it useful for the consumer, which again could be anybody. Let's just assume it's anybody. I feel like it's important to show business as usual, meaning things are healthy. The system's responding. Life is good.

We spent a lot of time in here about the doomsday scenarios or if a node goes down in some far-off land do we need to report on that? I don't think so. However, I think you're right. I think there are maybe some bigger buckets that you could put down and say that, "Should something happen in this category, we would make a statement on it," something versus nothing right now.

Because you're right. The RSOs have made statements recently. It's taken a while for the group to do that but they have. I think maybe, going forward, or should this go forward, maybe identifying a couple of those buckets of, "This might deem some sort of statement." Maybe it's not a disclosure as to exactly what happened. I don't know what

level of detail is in it but it's a statement, saying that—much like the RSOs have done with different things.

I think the big one that concerns me is route attacks, BGP attacks. I think that would be something that, if it occurred, I think we'd all be standing on the table, making statements about that. Availability, I think that's a tough one. Just because we are so dispersed all over the globe, I think that becomes harder and harder. But maybe we could create some minimum bar. I'm not sure. I just have a hard time thinking through that one.

But I think, as I talk through this with myself and listen to the conversation, I think it's easy to identify a couple of big buckets that the doomsday scenarios would fall into. Maybe you don't iterate through all the doomsday scenarios but you say, "In this bucket, if something were to happen, the governance structure would be expected to make some sort of statement that was approved by the RSOs," or something to that effect. That was my train of thought.

KEN RENARD: Wes?

WES HARDAKER: Thanks. Brad, I actually liked your notion of everything's good. The trick is when do you change that green bubble to something else? That's, I think, what we're trying to define, to a large extent. But really, I was piping up to say I agree with you, absolutely.

I don't think we've got to worry about the privacy side, in part because the rest of the world is already trying to think about that. When do you use TLS to an authoritative server is being handled in the IETF. There's even documents that talk about the root is the least needed to do TLS to because with QNAME minimization and everything else, your queries are not important.

The important thing is that, I think for that, I would wait for the rest of the world to tell us what they thought was private or not before we battle it here. The rest of the problems are more important for us.

KEN RENARD: Thanks. If we can go to the next slide. Oh, sorry. Duane?

DUANE WESSELS: I feel like some of these things, in particular the data integrity and the availability, are already well-covered by the metrics work that we've already done. We've already got thresholds. We've already got reporting requirements in there. And I'm not sure those rise to the level of what this work party is supposed to be about, which is security incident disclosures.

KEN RENARD: Right. So the metrics are going to look at availability. But at what point? 047 metrics were measuring individual RSOs as well as the RSS. So those have thresholds. Is there a different threshold for making a public statement?

DUANE WESSELS: I think those metrics reports are intended to be public. I think that's the whole point of it. So anyway.

KEN RENARD: Sure.

DUANE WESSELS: Maybe there's some overlap here or maybe we don't want the overlap. I don't know.

KEN RENARD: When does it turn from, "Metrics are above or below this threshold set in 047," versus, "When is this a material effect on the Root Server System?" I don't think they're the same. I think that you can fall below the thresholds of 047 and still have a very well-functioning Root Server System.

DUANE WESSELS: I don't think we want to reopen 047 at this time. But I guess I don't necessarily agree.

KEN RENARD: Okay. Maybe I'm mistaken on, when we started those numbers, where they came from.

DUANE WESSELS: At least we just need to be aware that we've already talked about some of this. And if we need to go back and fix it, fine, or refer to it, okay. But I feel like there's some overlap.

KEN RENARD: Yeah. I don't think there's any intention of revisiting those. That was really well-laid-out. But more along the lines of when should this be considered a security incident or an availability incident, whether it's a government or a regulatory body, should be notified. What do we owe to that community to raise the flag a little bit higher than just the published metrics?

Go to the next slide. We should probably move on pretty quickly to the 000 discussion. This is just talking about what roles could the RSS governance structure take? The first one is sharing data among the RSOs. Certainly, lessons learned are good things. We've been doing this, certainly, in different communities. But does the government structure have a role in facilitating this communication, not necessarily requiring it? Or is this something that just stays in root ops? Should the governance structure have a role?

And then, for anything that would be shared publicly, I think we've discussed this already today. Just an aggregation point for RSO disclosure. If anybody wants to look at events, even historical, of the RSS, here's where you can go to look at individual pieces. Or just after-action reports. Maybe it's a route attack. Maybe if the root server community puts together documentation of what happened—how it was mitigated and what happened.

Regardless of who authors those reports, the governance structure could just be a place to house and aggregate that information. So it would be a one-stop shop for anybody interested in the security of the Root Server System. Looking for any thoughts on those topics. If not, we can move on to 000.

All right. Yeah. Just additional things. If this work party were to continue, there's some very easy things that are out of scope. We could even continue this work party, thinking only of things discussed or disclosed among RSOs for those lessons learned versus going public. Again, I think the public part is something we, in some sense, owe to the Internet and will help us provide that transparency and continue the trust.

With that, thank you very much for the discussion. It's been helpful to me. I'm going to put this together in at least a summary document after the RSSAC and then maybe update the SOW for further discussion on that to see whether we go forward with this as a work party. Thanks. With that, Andrew.

ANDREW MCCONACHIE: Thanks, Ken. Can I share? Yeah. Perfect. You should be seeing my screen now. Good. My name is Andrew McConachie. I support the RSSAC. This next 30 minutes will be about RSSAC000v7, which is the seventh version of RSSAC's operational procedures. This is the third time we've had dedicated sessions to talk about this document. I think we're pretty stable on the text.

The purpose of the next 30 minutes is to just review. There's a couple nitty-gritty questions. But mostly, because there are such big changes coming in this version, with voting, and quorum, and elections, and stuff like that, I think it's good for everyone to just review this and really make sure that people don't have any questions about the text in the document. So I'm just going to walk through it a little bit. We'll just walk through it starting with quorum and voting. And then, if people have questions about the text or anything really, please just raise your hands.

Ozan, can you post the link in the chat so that people can look at this on their own computers? It might be a little bit difficult to see in Zoom. But starting off in quorum, so 1.4.1. This just clarifies what quorum is for an RSSAC meeting. It's basically a majority of Root Server Operators plus one. Then there was a bunch of confusion around precisely whether it required both or one of the representatives. We've just said, "either the primary or alternate representative is present." So that's hopefully clear now—more clear than the older text.

Moving on to voting. Again, here we tried to clarify that it's really the primary representative's role to vote but sometimes they're not available to vote. So then their alternative representative can vote. So we've used this language, "primary representative of each root server operator, or the alternative representative if the primary is unavailable." So hopefully that language is a bit more clear than there was in version six.

Then moving on to elections, this is where things get a bit more interesting. Previously, the RSSAC had definitions of majority for

voting in elections and really didn't separate out different types of voting actions. So it was just talking about voting. Then there was text about voting for the chair. There wasn't really anything specifically about publication of documents or motions. So in this section, we've broken out three different kinds of things that can be voted upon. There's elections, there's publication of documents, and then there's motions.

For elections, the RSSAC will adopt a ranked choice voting system. And then for publication of documents and motions, of course, those will stay at majority levels. So for publication of documents, it's going to require a supermajority of root server operators, which is independent of quorum. So if there's a vote within an RSSAC meeting, it's still going to require 75% of RSOs—not 75% of RSOs present but 75% of RSOs. Then for motions, I believe that's just a simple majority. Any questions so far before we dive into the ranked choice voting? Okay. Duane, go ahead.

DUANE WESSELS: Thanks, Andrew. I think our last meeting on this ended in this part, maybe, where one of my comments on the document is that we now have different ways of voting on things—publications, motions, and whatnot—and they have different bars. The bars are different heights. We didn't resolve that yet. Is that correct? I made a silly example where you could pass a motion that does something drastic and stupid to RSSAC, like disband it. You had a better example. But we didn't resolve that issue yet. Is that correct?

ANDREW MCCONACHIE: We did not resolve that. That's true. Yeah. You're specifically talking about motions. So voting on motions only—

DUANE WESSELS: That only requires a regular majority.

ANDREW MCCONACHIE: The issue is that there's no real defining of what a motion is. So you could have a really impactful motion that only requires a simple majority.

DUANE WESSELS: I just wanted to make sure I didn't miss something.

ANDREW MCCONACHIE: You did not miss anything.

DUANE WESSELS: Okay. We can come back to it if you want but I just wanted to make sure I didn't miss any changes there.

ANDREW MCCONACHIE: We can deal with it now because that was one of the questions I had. Your comment is still there. Do we want to put some definition around what a motion can be? That's one way to address it. Or another way to address it would be do we want to raise the bar for how difficult it is to

pass a motion? Or do we want to just enumerate what can be a possible motion? If it's just approving minutes, then we could just say that.

WES HARDAKER: It seems to me like somebody has solved this problem before us. Certainly, there are levels of motions. And certainly, passing minutes should probably not take a 75% supermajority. But certainly, modifying the operational procedures probably should. So I'm not really sure how to codify that really perfectly into a sentence. But I would think something along the lines of anything that changes how RSSAC operates, or its membership, or something like that requires the larger one. But I'd be tempted to do a little research to see. I don't think Robert's Rules of Order will do it, unfortunately, but I'm not sure.

ANDREW MCCONACHIE: We do have a separate section of publications of documents. So that wouldn't be included in the motions. So any alternation to and RSSAC document would fall under 1.4.2.2. and not 1.4.2.3. But there could be something else. Again, I am having a hard time coming up with a decent example. Robert, go ahead.

ROBERT CAROLINA: Thanks very much. Wes is absolutely right. This is a problem that has been solved routinely in a lot of different places. Just to give you an example, as Wes mentioned, Robert's Rules of Order catalogs, I think it's 98 different types of motion. I doubt you're going to want to

explore all of them since many of them are pretty obscure. But there are some basic principles could draw out of it if you wanted to look into this further.

For example, the basic idea is that a deliberative body can address the substance of their debate by making motions and taking votes on a majority basis. Where parliamentary rules—whether it's Robert's, or British Parliament, or anything else—tend to vary would be motions that threaten the integrity of the process of the body or moves that threaten the constitutional stability of the body typically require some supermajority. Roberts tends to default to two-thirds for things like a motion to suspend the rules or other types of things.

Then there are other types of decisions like the one you're looking at now, which is changing the actual rulebook, which require something else again, certain types of notice. So there is a way to draw clear principles between the types of decision-making that requires different levels of approval. And I would suggest that if you want to look into that further, there are sources of help that can be provided. I wouldn't want people to reinvent a wheel that's already pretty round and available.

ANDREW MCCONACHIE: Yeah, definitely. I don't want to go about reinventing wheels if it's not necessary. I did actually buy Robert's Rules of Order when I started working on this document.

ROBERT CAROLINA: That was your first mistake.

ANDREW MCCONACHIE: That was a big mistake. Yeah. How should we move forward?

KEN RENARD: Question, Andrew.

ANDREW MCCONACHIE: Yeah. Go ahead, Ken.

KEN RENARD: In some sense, the chair presiding over the meeting, could they be the one that decides whether this is a motion or something at some higher level of threshold—not that it solves all 98 cases but it gives a focus. On the other side of that, does that give too much power to the chair? Discuss.

ANDREW MCCONACHIE: Basically, put two options in motions—one for supermajority, one for majority—and then the chair would decide which bucket the motion fits in?

KEN RENARD: I think Liman has a great answer.

LARS LIMAN: I would be very uncomfortable with that because when you go into a process as a member of the committee, you must be able to predict how things are supposed to be treated. It can't be an arbitrary decision by the chair at that moment. So I would much rather see types of motions that are listed that need supermajority. And if there are typical clauses, as Robert suggested, maybe we can look at that and see if we can find clauses that we list there. But you cannot have a process that changes underway in the committee. It must be possible to deduct from reading the documents how things are supposed to be treated. Thanks.

ANDREW MCCONACHIE: Was there a hand up? Robert, go ahead. I'm sorry.

ROBERT CAROLINA: Yes. Just following up on this. I think that both points are well made. Interestingly, the answer in many systems of parliamentary law is that both are used simultaneously in the sense that you first draw principles between what is significant and needs to be protected procedurally and then you have what is not. Usually, you have to define what's very, very important.

Typically, what would happen in a deliberative assembly is that the person chairing the meeting makes the initial decision about how to classify the motion. But then there's a different motion, which would be to dispute the ruling of the chair.

One of the ideas of a parliamentary system or of a deliberative body is that the body itself is in control of the rule set. So if a chair's decision is not taken well, then someone typically would say, "Let's put that to a vote." And that vote, believe it or not, actually goes off on a majority basis.

So the classification argument can be put to the body. You need a central process point to initially decide, "How is this going to be handled?" people might even ask, "Madam Chair or Mr. Chair, how do you propose to handle this?" "I think this is a simple majority vote." "Well, I'm not sure it is, really." Or, "I think this is a two-thirds vote." "I'm not sure that it is, really." So we should test that.

Like I say, for me, it's an interesting conversation because I've spent time as a parliamentarian, advising bodies on how to do this sort of thing. But I would assure you that there are ways to solve this. There are ways to get through it—to balance the responsibilities accordingly.

KEN RENARD: Really briefly, the security geek in me sees so many holes in that.

ROBERT CAROLINA: The lawyer geek in me agrees.

ANDREW MCCONACHIE: Should I try to come up with some classes? I can go back to Robert's Rules of Order, and suffer through a few tens of pages of it, and

determine some good classes of motions that might be applicable to the RSSAC if that would be useful.

ROBERT CAROLINA: Very possibly. I'd be very happy to have a chat with you offline and make suggestions or give feedback if that would be of assistance. But I leave that entirely at your discretion.

ANDREW MCCONACHIE: I would be very open to that assistance.

ROBERT CAROLINA: I'm here all week.

ANDREW MCCONACHIE: Okay. Cool.

UNIDENTIFIED MALE: I would like to record that action item as, "Andrew will consult Robert and Robert's Rules of Order."

ANDREW MCCONACHIE: Is it a motion and should we vote on it?

ROBERT CAROLINA: See, now you're getting it. There you go.

ANDREW MCCONACHIE: So I think we've ran out of motion on motions. Any other questions on this section, section 1.4.2. Doesn't seem like it. All right. So I'm going to skip on down—we've got 15 minutes left—and talk about the electoral system.

I mentioned that the RSSAC will be transitioning to ranked choice voting. The open question was still, "Is this going to be applied to all outgoing liaisons or just the ones that are explicitly called out in the operational procedures?" I think in the last meeting, we didn't finish talking about this but we were edging towards just do it for all outgoing liaisons. Should I just record that as the way to go?

DUANE WESSELS: Yeah. I think my recollection from that meeting is we decided that was the simple thing to do and that's what we should do.

ANDREW MCCONACHIE: So I will record that decision. That gives us 13 minutes to talk about ranked choice voting. One thing I will say is staff looked into software systems that can do this and we found one. ICANN has a contract with a company called BigPulse, which has a whole bunch of different voting systems. So there was a bit of concern that we may not be able to figure out how this would functionally work. But I think we're pretty good on that now. So really, it comes down to how the RSSAC wants to design their electoral system and we don't have to worry too much about software.

BRAD VERD: For the record, can you just quickly describe ranked choice voting?

ANDREW MCCONACHIE: Is that a challenge? I know. No. That's fair. I'll use the example here. Basically, it's when you vote, instead of just voting for one candidate, you would vote for multiple candidates and you would rank them by preference. And you are doing a single ballot. As a voter, you would say, "Candidate A is my second choice, Candidate B is my first choice, and Candidate C is my third choice."

You have multiple tallies if necessary. When you tally the votes, you take everyone's first choice. And then, if someone doesn't pass a threshold—so, for example, you have multiple candidates. You don't just have two candidates. You have multiple candidates. You need to eliminate the candidate that has received the fewest number of votes. And you do that. Then you take the votes that went to that candidate, and you look at the voters and their ballots, and you determine what their second choice was. Then you apply that.

What it does is it makes sure that … The problem that RSSAC was having was when you have multiple candidates, you can't ever really achieve a majority because you're just splitting the votes across candidates. So this method just makes sure that you can get down to two candidates and still count everyone's votes by their preferences.

And there are some examples in the document. And there are some tricky—because the RSSAC will often have abstentions. That's something special for the RSSAC. The RSSAC needs to be able to deal with at least one or maybe two extensions. Also, I think voting

systems, in general, just break down as the number of candidates approaches the number of voters. The RSSAC only has 12 voters and can have three or four candidates. I've been looking at this like a math problem. So that's how I've been approaching it.

One outstanding question in this section was ranked choice voting doesn't completely eliminate the possibility of having a tie. So you can still have a situation where you have four candidates and the bottom to candidates have the same number of votes. What do you do in that situation? I think on our last call, we came up with two solutions. One of them is to vote again and the other solution is to just randomly eliminate one of them. How do people feel about those two? Do people have a preference to randomly eliminating the candidate with the least number of votes or holding another election? Wes, go ahead.

WES HARDAKER: Yeah. I hadn't considered that particular nasty problem. I think voting again makes more sense. Especially with only 12, random elimination could actually seal the deal for one of the candidates. I think that the important thing for considering ranked choice voting is the impacts on who it impacts. So for the voter, if they don't care about the algorithm much, they just need to know that they need to rank everything. Everybody knows how to do that. It's very simple.

If you do care about the algorithm because you're a geek like pretty much everybody in the room, you begin to understand the benefits to it. And one of the biggest touted benefits is it actually enables people to vote for who they want more as opposed to something like a two-

party system, which much of the world falls into, where you feel like if you don't vote for one of the two parties, you're going to be giving away your vote.

So it gets back to a system where I might pick to vote for somebody who's in the minority because I know that when they don't pass—because it's unlikely they're going to pass—my next vote will go to my next best candidate, which might be something out of the two-party system. So it tends to spread votes out. Or that's one of the goals. I actually haven't seen evidence that that actually happens.

DUANE WESSELS: I was going to disagree with Wes because I think random is better because it's more likely to produce an actual result and more time-efficient than re-voting.

WES HARDAKER: Yeah. It's hard. But there's still the case where you get down to two and they are matched. That feels weird to do a random vote that actually results. Then there's the case where you have one person that has, we'll say, six votes and the two people underneath have three. So you eliminate one. It makes more sense to eliminate both. If you actually have enough slots up above, you just drop both of the bottom two. But you get into the small cases and it gets tricky.

DUANE WESSELS: Yeah. It would be good to know what ICANN's software vendor does in this case, too—how they handle ties and things like that.

LARS LIMAN: Two comments. Whether you go for A or B depends on whether you see as getting an outcome is more important than the right of the individual that is the candidate. So I don't really want to have a preference on that. But I think that this is a problem, again, that has to be solved somewhere else. There are so many committees just within ICANN and many of them are small. So have we looked at how this is solved in other places in ICANN? Thanks.

ANDREW MCCONACHIE: I looked around to see who else was using ranked choice voting. I didn't find any instances of it within ICANN. I know DNS-OARC uses it. I looked at small municipalities in the United States that use it. That would be even smaller than any municipality, I think. You only really have ties when you don't have that many voters. So it's a particular problem for a small committee to do it because you only have 12 voters. Even if you have a small city with 5,000 people in it or something, you're very rarely going to have a tie.

WES HARDAKER: What does SSAC do? I actually don't know.

ANDREW MCCONACHIE: For what?

WES HARDAKER: For electing a chair.

ANDREW MCCONACHIE: For electing their chair? They generally have a very open discussion and they do vote. But again, it's just majority. It's just whoever gets the most votes wins.

WES HARDAKER: So they don't have a 50% requirement?

ANDREW MCCONACHIE: No.

LARS LIMAN: Whoever gets a majority is a different thing from whoever gets the most votes.

ANDREW MCCONACHIE: That's true. But as long as I've been supporting the SSAC, I've only seen one chair election and then it was just really not contentious. So I couldn't really say.

WES HARDAKER: We could do something crazy and say we revote. If no change is made in the numbers, then you do random.

ANDREW MCCONACHIE: Yeah. I guess my concern with revoting is you get in this endless cycle of revoting. So maybe you do like you suggest. You revote once and then after that it's random. You get one do-over.

KEN RENARD: Andrew?

ANDREW MCCONACHIE: Go ahead, Ken.

KEN RENARD: One of the purposes of this session was to have the broader RSSAC really look at this document because there was a couple of us on the call. So quorum, and voting minimums, and dealing with absentee or abstentions are in this document.

Just encourage everyone to look at that and make comments in the document. It doesn't sound like we're going to be voting on this anytime soon. But one of the goals was to have this new voting procedure available for the chair election, just as a random goal—not absolutely required but let's shoot for that. So if we can have everyone take a look at that, especially some of the quorum, and minimum number of voters, and abstentions, and provide comment in the document. Thanks.

WES HARDAKER: So if I can reinterpret what you just said, your goal is to have this done by the November-ish—I don't remember the exact timeline—election.

KEN RENARD: Yes. November-ish.

ANDREW MCCONACHIE: So I guess that would mean we'd have to approve this document in the November RSSAC meeting.

WES HARDAKER: Or October.

KEN RENARD: Yeah. And again, it's a goal. It's not a requirement.

WES HARDAKER: This is where I'm staring at Ozan. I think nominations go until end of November and I think the election's in December. You guys recently got elected. You should know.

OZAN SAHIN: In terms of the timeline, we are planning on starting the call for nominations for the RSSAC chair elections towards the end of October. So that's going to be a 30-day call for nominations period. And the actual elections will take place during the RSSAC December meeting.

WES HARDAKER: Thank you, Ozan. Perfect.

ANDREW MCCONACHIE: Then, yeah. I think we'd need this approved on the November meeting. So this should be approved one month before we actually do the election.

So wrapping up, I have an action item to work with Robert on some language on motions and to add some text on what to do in the case of a tie. We're going to go with one revote and then after that it's random. Then I think we're done.

KEN RENARD: All right. Thanks. Do we have another RSSAC work session this week or are they Caucus Work Party meetings?

OZAN SAHIN: Later today, we will have two RSS GWG work sessions in this room. Tomorrow, there will be other RSSAC and RSSAC Caucus work sessions.

KEN RENARD: Thanks. And with whatever power I have, I say adjourned.

[END OF TRANSCRIPTION]