JEFF OSBORN: Good Is this thing on? Hello? Can you hear me? Okay. Okay. Excellent. We, do not have a wireless mic. So Ozan wants to try something. I think it's gonna be really kinda interesting. So we'll see whether it works. We're gonna line up and introduce ourselves at the mic and go sit back down. It's gonna be like at the old Baptist church where you went and got communion one at a time, and that was usually a cluster thing. So let's see how this goes. Dave, you wanna start?

DAVE LAWRENCE: Hi. Dave Lawrence Tale, Salesforce.

PAUL HOFFMAN: Paul Hoffman, I can, who gets up at mic quickly.

TERRY MANDERSON: Terry Manderson, ICANN, IMRS.

KEN RENARD: Ken Renard, Army Research Lab.

ANDREW MCCONACHIE: Andrew McConachie, ICANN, Policy support staff.

SHUMON HUQUE: Shumon Huque, Salesforce.

DUANE WESSELS: Duane Wessels from Verisign.

HIRO HOTTA: Hiro Hotta, WIDE and JPRS.

KAZUHIRO KITAMURA: Kazuhiro Kitamura, JPRS.

SHINTA SATO: Shinta Sato, JPRS.

ROBERT STORY: Robert Story, USC ISI.

ROB CAROLINA: Rob Carolina, ISC.

YUJI SEKIYA: Yuji Sekiya, WIDE.

TIM GLADDING: Tim Gladding, Verisign.

KAZUNORI FUJIWARA: Kazunori Fujiwara, JPRS.

ALEJANDRO ACOSTA: Alejandro Acosta, LACNIC.

HAFIZ FAROOQ: Hafiz Farooq, RSSAC Caucus.

SHAILESH GUPTA: Shailesh Gupta, RSSAC Caucus.

RAY BELLIS: Ray Bellis, ISC.

JEFF OSBORN: Thank you all very much for that. That was way more orderly than I imagined. I thought we could spend 20 minutes disengaging from that, but it worked pretty well. So thank you. That's a call to order. We're taking a look at the agenda, it is it is fairly basic. We have some, administrative issues. And primarily, we're gonna hear from, where you're gonna have that work party meeting with, Robert Story. So, moving on, does anybody have an agenda item that they're thing here. I'll give you a minute to take a look.

Okay. So heading back to the beginning, We would like to, introduce the RSSAC caucus members who are new since the last time we did this. And my Macintosh,reconfigured itself and added an operating system in the our son, none of my notes are where they should be. Oh, but here they are. So since June 2023, I'd like to welcome Brad Harris from NASA, Joe Hayes from DISA, Jose Nunez-Zapata from NASA and Peter Thomason who is with Oh, .de nic? .desec, okay. Also a member of SSAC.

Great. So welcome to the new folks. And that gets us to work parties and, work products. Now this is this is an interesting section where we put out a survey that many of you were able to fill in, and there's a free form section at the end. And so I wanna state that we we wanted to share the free form requests for things to study were without calling this anything like a definitive list and here's what we're picking from rather those of you who finished the survey and did put in something free form these are the pieces that came out.

So, it came in 4 sort of 4 chunks. And the first was pretty obvious. DNS, service stability, the root KSK rollover and TLD studies, which is just about what you'd expect. The second was a little more, specific I'll I'll leave it to read there or read it out loud. I'll read it out loud.

1, Many new attack that related to DNS are discovered regularly. Some more discussions on these need to take place.
2, DNS threat modeling, 3, domain generation algorithms,

For MITRE attack framework and it's mapping to the DNS threats. 5 new naming systems for decentralized DNS, handshake domains, HNS, Ethereum name services and web 3 domains. So somebody has an interest in security and attack protection.

The 3rd primary writing was the impact of hyperlocal root on service levels and responsiveness. Variance of SOA serial coherency within anycast in and across letters and the use of zoneMD and the 4th is the root zone was of significant importance in the WGIG process leading up to the WSIS forum from 2003 to 2005.

Think it's important to honestly reflect on the discussions and how the root is now managed differently or not after 20 years have passed. Well, the root zone is managed well. Thank you. It is important to demonstrate that as well as proactively identify ways in which it can be improved, all of which become their own contributions to the upcoming WSIS+20 review as well as the UN's SoTF.

So I throw those up for your consideration. I believe we have no new work groups being launched in the short term or anything on deck. This is a set of those. There have been other things brought up and there are certainly other channels for for proposing workgroups. But we felt it was the responsible thing to read them because you guys went to the trouble of writing them. Let's see. There's a hand from Shalish. So Shailesh. Old hand? Oh, no problem. Okay. You've got the rest. Anyone else? Okay. Our next agenda item..

KEN RENARD: This is Ken. Probably the idea of of discussing or having this on the, list is bring up here in the meeting is if anybody that did suggest any of these topics wanted to discuss them more or you know, make a pitch or try and get folks interested. Certainly, there are some things in this list that are really outside of the root server system, that, if you don't, they're very interesting not don't necessarily fit into the the caucus work. But, if anybody did, did suggest these and wants to talk about them further, please feel free to say something here, if you'd like.

And there was one Duane, you had brought up a potential work item want to talk real briefly about that?

DUANE WESSELS: Thanks Ken. This is Duane Wessels. So, what you're referring to is something that, we talked about months ago, which is to maybe develop a policy around renumbering of root servers or or change, changing a root server's IP addresses. I think that would be a very interesting and fruitful work party with within RSSAC and oh, something I'm definitely interested in and would volunteer to work on.

JEFF OSBORN: Definitely an interesting one, Duane. I believe Robert is approaching the mic.

ROBERT STORY: Yeah, Robert Story USC ISI, I would also volunteer to work on that with Duane, obviously, be as very interested in renumbering and policy and developing how that could work in the future. Thank you.

JEFF OSBORN: Thank you. Anyone else, remote, in the room. Okay. Can you get, oh.. Hafiz?.

HAFIZ FAROOQ: Hafiz Farooq. Thank you very much. Actually, I think the security, topics, I'm, recommended by me. And, they were recommended at the time when there was no working party, or I was not involved in this the incident reporting. So maybe once we have incident reporting practice in place, then maybe at a later stage, I mean, security, these topics which I'm recommending. They can discuss later on. Not at not at this point in time because they might not be related directly with the with the RSSAC. But once we have more incidents coming in future, reports coming in the future. Maybe we have more visibility that if we wanna do some working party later. Thank you.

JEFF OSBORN: Thanks, Hafiz. That makes sense. Anyone else? Can we flip to the agenda? And I believe that gets us to item 5. Robert Story, sir? If you wanna sit there, this mic really works, and I'm not plugged into anything so you can have all.

ROBERT STORY: So, you are going to start with the normal welcome and roll call? Are we gonna skip that? Cause everybody already introduced themselves at the mic.

OZAN SAHIN: Yes, I think better to skip that, but, maybe just an announcement that we are now transitioning to the, 10th meeting of the RSS Security Incident Reporting work party. Over to you, Robert.

ROBERT STORY: Okay. So the second item on the agenda, and as always, this agenda bashing anybody would like to throw out additional items for the agenda besides working, looking at the working party documents, nobody's running towards a mic. So, I guess we'll just, jump right into it. Yeah, you can switch over and I haven't looked for other people's changes. I'm made some changes. I know Andrew did as well, so I guess we'll just scroll down until we find changes and decide if we want to accept, talk about. Anything? Okay.

PAUL HOFFMAN: So, this is Paul Hoffman. For those of us who haven't been active in the work party, I sort of expected a summary first, rather than jumping in because I also haven't read the document and,
Or I'm sorry. I looked at the document enough to notice that there were lots of paragraphs that were fine and lots of paragraphs that had been heavily scribbled out. Can you give a summary of where you think the work party is and where the document is now for those of us who are, This is in the IETF. It's okay to have come to a meeting without having read the draft.
Can you give us a summary of just help help help how you think it's going so that we understand what's going on.

ROBERT STORY: Okay. Definitely I'm not prepared for that. I will make a note for future meetings to, to to to try and prepare a summary. You're right. I'm excited.

RAY BELLIS: Sorry. Yeah. Ray Bellis. I mean, I'll top, we've not been prepared. Actually, I think this is kind of unprecedented that we're actually having a working group session or work party meeting, during an IETF hosted RSSAC caucus meetings, things it's never happened before. It's caught me out by surprise. Was not expecting this. I was expecting just a 500 summary on what part has been up to.

ROBERT STORY: Okay. I was not in charge of the scheduling. I was just told to be here and that I was chairing a meeting.

OZAN SAHIN: Yes, thanks for this question. This is Ozan speaking. So I guess, this is the, this is what the, RSSAC Admin Committee wanted to

do this time because of the.. We had another RSSAC caucus meeting 10 days ago, I guess. So, during the ICANN annual general meeting, at ICANN78. So, we thought we thought we would we would cover the general RSSAC caucus agenda during this meeting and give most of the time, to the Work party meeting, during the RSSAC Caucus meeting at IETF. Because of the very small, you know, because of there's just only 10 days between these two meetings.

ROBERT STORY: Okay. So,Paul

PAUL HOFFMAN: So, Robert, if you're not able or willing to give a summary, if somebody else in the work party is willing to, I, I understand that asking somebody to start riffing is, is a little bit unfair. But it still would be useful at least two of us in the room to be hearing it if somebody else who's been active in the work party willing to say where they think the work party is at.

ROBERT STORY: Yeah. I'm gonna get revenge on Ken and throw him under the bus this time.

PAUL HOFFMAN: We can do that at the same time.

KEN RENARD: So, thanks. This is, Ken Renard. So this gone even back I could even further. So this work party, and it came out of RSSAC58: Success criteria for the governance structure that says the governance structure should have mechanism to due security incident reporting. So, there's a lot of reasons for this. A lot of it based on transparency of the root server system and It's also been a target or a something mentioned by, several types of regulation that called out the root server system for the RSOs specifically to say they should do some sort of root server or some sort of security incident reporting.

So, transparency, that's a good thing. So that's kind of what we're trying to mainly focus on here. The statement of work, kinda talks to really just that. The Work Party thus far has been just spewing ideas onto page. So this document looks a little bit, you know, disorganized and and, you know, lots of things crossed out. So that's kind of where we are. We're try we're we're about to turn that corner towards trying to get to a couples couple specific targets of what we want this this document to look like. And, welcome thoughts opinions here on on where this should go. And

I think I'll let oh oh oh oh oh pull Robert under the bus group. He's out of anything else.

ROBERT STYORY: Yes. So today, we have been trying to go back and forth and try and and find the the right balance for how we define what qualifies as a security incident related to the availability confidentiality, integrity, of the root server system. And so we like a lot of committees and working parties where anybody can participate in we have trouble going too much into the details, and so we're trying to find that that balance of the right level of of information to to have in in the document. So, And that's I started and that section that I worked on later trying to to to pull back and come up with more general definitions. We've just to have talks many times about the decision about what might be qualifies a reportable incident is always gonna be subjective. So we can't give an exhaustive list of things that should be reported and so we're working on trying to find the rights definitions are examples as as guidelines. And so that I think is essentially where we're at now. So Andrew has, got a new section here to try and get some terminology. We were using various terms all over the place, and so we're gonna try and be more concise and use, find the terms we're gonna use and use them consistently.

ANDREW MCCONACHIE: Yeah. Let me just kind of explain what I did here. So I was tasked creating a terminology section. I I started off by just saying, you know, what we typically say in a lot of our sec publications, which is we're gonna use everything from our RSSAC026v2, which is the RSSAC terminology document. And RFC8499 know, recognizing that there's a basing, progress for that. And then I, thought of some definitions. I I, Robert Carolina had put this long footnote, on page 4, and I kinda I I stole that footnote and turned it into a definition. Same text, different place. The other the the CIA, you know, the, confidentiality availability and integrity, I stole from NIST. And then I tried to add like, another sentence that kind of qualifies it for the RSS. And was are there any other definitions there? Scroll down a little bit more. Ah, yeah. Reporting is another one I just stole from NIST and then security incident another one. I stole a lot from NIST, basically, is what I'm saying. Because it's always, in in my view, it's always best to steal terminology than create your own. I don't know if we wanna review those those terms or, add new terms or think about how I modified the CIA ones to be more for the RSS. If you wanna do that now, if do that later.

ROBERT STORY: So, as I said, I haven't looked at other people's changes, I think, they're definitely, will probably be some comments on the the definitions of for example, security incidents, I guess is one thing we've had discussions about terminology for perhaps security incidents that might not the reportable you know, where the threshold is for what something that needs to be report it. So but I would wanna think some more before making comments or suggesting text. Anyway here, read the documents or is reading it that would like to comment on these sections? Whether these definitions? And I see someone, Hafiz

HAFIZ FAROOQ: If you can go to the reporting definition, I mean, they start the definition by saying this is the final phase of if you can scroll, this one. So it's a final phase of computer network forensic process. So but our incident reporting, we are talking about initial report and the, the final report both. So maybe we need to revise this definition. I know it's coming from NIST we need to put something which is in line with with you're looking Thank you

ROBERT STORY: Right. And we've also talked about different, levels of reporting, in addition to security incident reporting, we've also possibly talked about transparency, reports. So, anybody that wants to contribute, wording to the definitions, feel free. Unless anybody has any comments right now, then we'll just keep working through the documents.

We can scroll down and look for more changes. So this is a section that that I added here under, reportable incidents. Where we're trying to pull back from, to find a level of definitions into coming up with sort of catastrophic things that are without a doubt something that would need to be reported and so just from a document that it had put together earlier and some comments in the the past meeting, through together this, quick non exhausted list of some accidents that would likely be reportable.

So, in our RSSAC002, in the measurements, we talked about the K of n parallel availability, I'm sorry. RSSAC047. Where The definition was that, having 8 of RSOs being available would give us the 5 9 availability. So any outage of 5 or more RSOs would definitely qualify as as it's something that we should take notice of it and and report. I also added the unavailability due to, BGP hijack or RPKI failures. I don't know if that could be a sub bullet under a complete out outage or if a hijack or RPKI failures of a a

single RSO is reportable, or would it be the same Yeah. It affected 5 or more level thresholds. So, open to to comments on that. In the integrity section, of, Basically, at integrity, we don't have a percentage other than 100% in in integrity. So any RSOs, serving incorrect data either by modifying the zone, filtering the zone, playing with DNS signatures, is, would obviously be bad And then we have had some discussions on things that are outside of the control of the RSOs. And but I feel like some things such as getting invalid zone data from IANA. That's not our fault. But it still seems like something that we should, that we should talk about. And I see I have triggered comments. Duane, go ahead.

DUANE WESSELS: Thanks. This is Duane Wessels. My my comment is actually about the availability bullets. A So I the second one in particular, the the unavailability of of one of our associated BGP hijack, or RPKI failure. I think To me, the interesting part of that bullet is the BGP hijack or RPKI failure, not, not so much that it resulted in unavailability. Because I I there's a very stark difference between. The second bullet in the first one, which requires basically unavailability of 5 RSOs, for, for some reason. And then the second one talks about unavailability of one RSO for this particular reason. So I don't I don't I don't don't understand why it's different. The interesting part about the second one is the RPKI stuff because that's like a high target point of attack or or or or something like that.

ROBERT STORY: Maybe, a hijacker would be somebody else taking control that could be, integrity

DUANE WESSELS: Whether or not have resulted in a unavailability or not, just the fact that a hijack occurred is maybe something that's the reporting is my point.

ROBERT STORY Okay. Hans Petter.

HANS PETTER HOLEN: Thanks. Can you hear me?

ROBERT STORY Yes.

HANS PETTER HOLEN: Excellent. Thank you. So I think both for integrity and availability, my question would be regarding to the time dimension. So I thought about it when you said serving incorrect data. No. I don't

think anybody would, on purpose serve incorrect data, but it could easily be that somebody serves old data. No, the question is how old hump to be before it's old. So I think here, both on the integrity part and also then on the availability. Complete outage, but for whole long, a 2nd, 5 minutes, 3 hours, 3 days? And then I also have the same comment, here on why our some causes for an availability causing us to report when the the general criteria is 5?
And I think with the scrutiny upon us, I think we need to be prepared through report on all availability of all the individual loans, but make it clear that, you know, unless it fits the threshold of 5, it's not, material. But I guess that gets into the the details once we start to discuss the timings here and also the the thresholds.

ROBERT STORY: Right? So, again, this is just a list I threw together, last night. So and, specifically to have something for people to throw stones at. So yes. Exactly. And, so we have several times talks about specifically not wanting to single out reporting, I think, for single RSOs, And, so defining the the the the thresholds or suggested thresholds. Again, this list is not meant to be exhaustive is as, part of the process. Think we had one more, well, Paul was had his hand up online right before you, Terry.

PAUL HOFFMAN: So this is Paul Hoffman. So my question really is who's going to do this reporting? If I'm, you know, t root, or p root, I guess, would be who I would be. And I have a full outage, I don't know if there's 4 others. I can't say, I'm sorry, I tripped the limit to 5. I can say I'm fully down. I can do a report either at the time or later. I know you have on that. But I have a concern with what is here and what you just said, which is, well, there's gonna be some thresholds for enough. Is the reporting going to be done by individual RSOs, or is there going to be somebody that says we're looking and now 5 are down, and therefore, we have to report. Those are very, very, very different scenarios.

ROBERT STORY: Yes. And we have talked about that.

PAUL HOFFMAN: Is it in the document that I just haven't seen?

ROBERT STORY: It it is in there. We've talked about the scope of work says that availability integrity, compromise that materially materially affects the RSS system as a whole. And so one RSO disappearing does

not materially affect. And so there's been some, questions about definitions that may be used that was in the, footnote that Rob Carolina had had about the governance structure and and saying that root ops as it currently exists now is a de facto governance structure, even though it's not actually governing its its collaborative, but but anytime that there's, a significant event. At an RSO. Generally, there will be collaboration and and is anybody else having a a similar issue and and we have a procedure in place where people can suggest, well, something happened, which may have then at their own RSO or or other events where we have can come to a consensus of where we wanna make a a statement from from root ops. And whether or not what happens when the governance structure is in in place whether or not there's an additional level or kind of there's a committee or something that is, subpart of the governance structure that takes over the role of root ops or whether it then there's multiple levels still remains to be seen.

PAUL HOFFMAN: But for now, it would be root ops reporting that it was significant. So your second bullet either is not significant or that operator, you know, like, when when I break my routers, which I have done. When that happens, would that report come from me be expected to come from me or come from Root ops going: Paul broke everything - I'm sorry: Paul broke all of his things. None of us even noticed.

ROBERT STORY: So I think that's still up in the air. Another discussion that we've had is are some incidents more significant or worthy of reporting than than others.

PAUL HOFFMAN: I'm really asking the who is is with without because once we have levels, we're gonna redo our RSSAC 47 and such like that. That's all fine but reports have to come from somebody, and either they come from just the individual at which point they can't do the 5 where it comes from just root ops or successor, but it might be about 1, or it can come from both. That that should be made clear.

ROBERT STORY: Right. And I I don't think we have clarity on that yet.

TERRY MANDERSON: Terry Manderson, ICANN. As an RSO, and this is only a a technical knits with the integrity, as an RSO, I don't receive the

zone file from IANA. I receive it from the root zone maintainer. Also look at, compromise of the root servers dot net zone, there is zone. They're just entries.

ROBERT STORY: Domain. I guess, for a better, better word.

ROBERT CAROLINA: For those who may be remote, there is a fist fight. And they're swinging. They're off.

ROBERT STORY: Rob, go ahead.

HANS PETTER HOLEN: It's not possible to catch any of that discussion remotely.

ROBERT CAROLINA: Yeah. A couple of a couple of quick things. That have kind of been touched on already. And just to highlight some of the things that were discussed in Hamburg, as well. One of the difficulties that I think earlier versions of this work product have had, and I think this is in the process of being remediated. Is the whole distinction between what is it that the the Collective operators, the root service system are going to report to the world as opposed to what is it that individual RSOs are going to work to that collective governance structure. Part of the context of this is that in the in the in the context of discussing and developing the next stage in the evolution of governance structure in our RSSAC058, Success criteria 8.1.a.1.1.1. Specifically was added to refer to including within the revised governance structure, a mechanism of both collecting up incident reports from RSOs in anticipation that we wanted the evolved governance structure to be a source, a repository of information and question answering for some has yet to be determined number of interested parties around the world who insist upon that reporting. And within the document, I think it's moving in this direction, but the just to highlights something that I keep saying and these things. Is that we need to be clear about what is it we think is reportable by the governance structure to the world, And now we've got sort of like two elements of that that have come up in 2 different meeting rooms. One is what is reportable in the sense of an incident in a very short time frame, and then what is reportable longer term as part of transparency reporting, which does not may or may not fall within the scope of.. I don't wanna make this document and [INDISCERNABLE]. Answer to everything. And then the second thing, which I think I don't know if it's going to be teased out in this document, But in order to do that first piece credibly, there

must be some understanding of the amount of reporting that each individual RSO was prepared to expose to the emergent government structure. I mean, we see that in the chart that's a little further down on this document, you know, socializing an incident, but what there is not a discussion of is what is the trigger for compelling is too strong a word but for explaining to an RSO that this is the moment when you really ought to be talking to the other RSOs about an incident so that the collective can then figure out what is or is not portalable

Now last thing, and then I'm out for a little while. And that is in a lot of the new language that I see here something that seems to recur in in documents is I think there's a temptation to jump far too quickly to what I would call a highly defined Operational standard. No more than 5 of these. No less than 6 those no more than x number of milliseconds, things that are easily measurable you know, with straight edge, a ruler of of of physics lab. I think that I think that the document would be stronger. Each time it starts to break into a discussion of that would be to begin with a clear statement of threshold in the sense of that would be understood by an external party who relies upon the system. And here, we start using words that engineers will inherently feel are mushy and meaningless, like material impact or significant or significant disruption of words like that, which are you can't measure in a physics lab, but if you get a hundred people together and point them at a set of facts, they will quickly converge on whether or not this happened in such and such a case. So they those types of terms can be much more useful.

I would I would urge the working party, and that's one of the reasons why I recently pledged to work with the working party. To focuse very clearly on setting those functional boundaries before jumping immediately to So, like, the the tools on the workbench of how we're going to measure what that means. I think that will help focus the discussion, and it will certainly help to reduce the tension in the minds of external parties like regulators who just really want to know that we've got this under control.

ROBERT STORY: So that's a constant battle that we're waging in in the documents is being too specific or too general. And so I I think it is excellent to have input from the nontechnical side, and this list here was was my first pass at an attempt to pull back and and list things

that are are so obvious that qualify as a material effect. And so maybe that where it needs to go in here in in a couple places. So like I said, it's it's here for for everybody to throw stones. And so please throw stones are even better, write some text. Ken, you look like you wanna say something?

KEN RENARD: So one of the things we've discussed in this past couple meetings is looking at, you know, is that there's confidential confidentiality integrity availability is just plastic security definition, know, confidentiality is something that doesn't is not a big of a deal to the root server system. But, you know, for each of these we've talked we can't define this. It's subjective Let's have some examples. So that's where this is. That's where this list comes from. Alright. I apologize I haven't gotten to it myself, but I think as far as a document, if we talked about availability here are some general criteria about how how availability is subjective. Oh, and here are just a few examples not not exhaustive. Things like that. So, I will probably add some text you know, to to introduce the availability stuff. If anybody else wants to at some just general text of what what what integrity means, what as far as what should be reportable. Again, describing its its subjective This is our Zen of what we mean for of what it means to be reportable incident for integrity invite folks to do that.

PAUL HOFFMAN: So I have a question for Ken now. Which is where you just mushed everything up. What happens with RSSAC047? That is what I've been hearing up till now. and again, I haven't been spading in the work party, but what I've been hearing up to now is the work party would possibly change some of the stuff from 0 47 and such but we can't make 047 mushier. So is the intention here, eventually, to be solid in a way where 047 could then help inform the root server operators and the community what's going on. Or is this now diverging an 047 will keep its fixed numbers, and this will be mushier and more advise.

KEN RENARD: Thanks, Paul. My my take on this is that 047 does a great job measuring the root server system. It keep keeps an eye on it. This is this is our standards, this is what we want. It can be used as a basis for security incident reporting or for, you know what we present to the outside world, but it's not you know, the 47 had that specific you know, measuring the system. This is more just security incidents.

ROBERT CAROLINA: Yeah. I'd take that a step further. I mean, 047, to my mind, functions while I'm I I was about to use a poorly chosen phrase that I don't wanna use. It is a very useful tool to measure expectations of how the RSS collectively is meant to be operating. And I take as gospel the word from all of my technical friends who say that it functions very well in that in that respect. I think that what we're dealing with here though, is a document that is meant to amongst other things provide comfort to those people in different parts of the world and in different administrations and bureaucracies who want to know that our small corner of critical infrastructure is being looked after properly and that if something of, let's call it significance happens. That we will be telling people about it. So to my way of thinking, that's the most important thing to get from this document as a statement clarity to the world that if something of significance happens, that we won't just sit on it. We will be telling people about it. We will be explaining to people, how it's being addressed. To accomplish that mission, I don't think that 047 on its own can possibly accomplish that. So we need something where broadly that that moves in that direction, which is one of the reasons why I keep talking about describing standards. And by the way, I'm I'm I'm going to wanna throw the gauntlet down every time I hear that somebody say subjective, a subjective decision or subjective standard. I'm I'm prepared for me to challenge that. Because the types of standards that I'm referring to, things like a material risk of interruption. That's not a subjective standard. That's an objective standard. It's a standard that measures something that can't be measured in a physics lab. It can't be measured on an oscilloscope. But it can be measured within the framework of reference of finding a group of people who are skilled and understand the system, and just pull the Did this present a material risk of disruption? Now if you find 90% of experts in a particular field converging on a single answer.

To that question, I submit to you that what we have is not a subjective standard. It's a objective standard that is merely designed to reflect what you might think of as accepted, accepted scientific opinion or accepted engineering in So, the and the reason that I'm preparing to challenge the use of the word objective is I think people use a as a way of depressing or reducing against the standard that's being adopted. There's a lot of value, There's a lot of value in words like material. There's a lot of value in words like reasonable. So, trust me. Don't throw them don't throw them away. They're they're some of our best friends.

ROBERT STORY: Yes, we have, remote hand from Hans. Go ahead.

HANS PETTER HOLEN: Yeah. Thanks. So I want to to build on what what Robert just said and, usually when you try to build incident reporting, the incident report procedure describes how you do incident, detection mitigation, root cause analysis and reporting in the end. Then you have another document describes your service level objectives, and then the definition of an incident is usually something that, violates or make sure that you don't so something happens and you don't reach your objectives. Right? So the the objectives don't belong in the procedure itself. RSSAC047 is kind of like ish an object, these objectives, but it actually says only advisory on metrics. So unless we commit to them, they're not really there.

So that's why we just moving into this discussion on now, what are really the thresholds because we have discussed how to measure them, but not exactly what we're committing to in all cases. And the then to make life even more difficult than some of the schools of incident also defines something that potentially can put you in a state that you can't provide service as an incident. And then what the expectations from the people around those is if you read, for instance, NIS2 directive, which we're not subject to, but there is an expectation that we would behave according to it, is that if something is happening of material, effect that, that, was was just mentioned, there is also a requirement for an early warning report within 48 hours and so on. So I think understanding the the requirements in all of these different legislations from all of different governments is exactly what are they looking for and how are we going to to deliver on that is is part of the challenge for this, for this work party, and I don't think we will add to the end of that, unless we have also agreed on the service level objectives, which you know, is probably outside the scope of this work party, but I think we need to crack that at some point as well.

TERRY MANDERSON: Terry Manderson. Hi, Ken. Firstly, I retract my comment about, rootservice.net. Oh, I had a brain fail. It's 2 AM. My body clock is wrecked. But I I would like to ask a question about confidentiality. And and why it actually applies here in terms of the DNS protocol when you're talking recursive resolver to authoritative server. When we're thinking that q name minimization is probably a a better answer, and anything that impacts the the query and response at that level, you're then talking about integrity. So but

I might have been for for a discussion on that. I'll I'll just It just doesn't gel with me. That's all.

KEN RENARD: This is Ken. So, I get confidentiality was put in there because it goes with the other 2 in every security textbook you have ever seen. So, yeah, the what's what are the confidentiality expectations in DNS. Pretty much zero. So, I think I don't know if I put these in what, what's something, anything that could potentially be an example. So, if we if we want, if we want to just completely get rid of confidentiality, that's fine. It was there more for completeness on the security side.

TERRY MANDERSON: So, in response, I think your your first statement was it's basically zero, I think put that in there. And say it's there, but we don't have that as a major concern that needs to be addressed in this document.

ROBERT STORY: I just wanted to to respond to that. I mean, there have been providers like Apple that are looking into privacy preserving for for DNS. Specifically to protect the the contents of the query, what information people are asking for. And you're right. Q name immunization does that between a resolver and the and the authoritative. But we can't tell the resolvers what to do. It's up to them to whether or not they're gonna implement or turn on like, Q name minimization. So I think it's still open for a discussion, just my two cents. And Hans again.

HANS PETTER HOLEN: Yeah. Thanks. So I I was one of those arguing for keeping us with the terminology and modifying the terminology when we spoke at the ICANN meeting. One of the, one of the puns on, on security incidents is, the CIA, as we've talked about, but of what it's not about the protocol, it's about the network system and applications. So it goes into the CIA of the NSAs. So, you know, if somebody breach confidentiality of your servers, well, sorry. Your, your service may be, maybe compromised. So I think, you know, stating that there are no confidentiality aspects of the protocol that shortcut off the system if, the admin credentials are not kept confidential, then we are in the trouble. So taking out confidentiality of the reporting framework, I think, we are not doing our self in service.

ROBERT STORY: Okay. Nobody else online. Ozan tells me that we have about 5 minutes left. So I guess we can continue just go into the

document and see if anybody else has made some changes. A lot of these changes are from the previous meetings. I think probably as the Admin Team we should maybe go through and try and clean some of this up, before the the next meeting so we are not scrolling through multiple times. So, unless anybody has anything, they particularly remember commenting on or that they, would wanted to discuss, then I don't know if we have 5 more minutes stuff to to talk about it, maybe we can move on. Well, yes. Right. Yeah. I'm I'm just, are we ready to wrap up this portion is what I'm saying. And I think we probably are. So yes.

OZAN SAHIN: Thank you, Robert. So in terms of the, timing of the next meeting, I guess we will, schedule the next call on 20th November, and I'll send out calendar invites for that meeting.

JEFF OSBORN: Thanks, Robert. An item that we added, I think, for the first time, in the caucus meeting is gonna be Ken, acting in the role of asking RSO. And this is something we've never really, I don't think had the time or the setup well to do. We seem to, like, run out of time and say “any questions?”, “sorry, gotta go”. And so, Ken had offered. There are no stupid questions Is there anything people on the Caucus who are not regular attendees to the RSO meeting or the RSSAC meetings would like to ask a root server operator. And if there are none, I'm gonna let this sit here silently and hang in the air for 5 minutes, and then we'll go to the next thing.

KEN RENARD: This is for you know, caucus members if if they that are not involved day to day with rso operations, if there are any questions, or even about specific work parties. We just talked about the other one, but, you know, we wanna make it easier for caucus members to join participate and, you know, share their expertise. So any questions?

SHUMON HUQUE: So I'll make a comment. It's not really a question on the confidentiality thing because I almost said it like, but then I didn't wanna open a can of worms, but I don't think it's really true completely say there are no confidentiality concerns with the root server infrastructure. Because, even if you, first of all, not all resolvers use Q Name minimization. So there is a confidentiality exposure for individual queries. Right? And that's number 1. Number 2, is even with Q NAME minimization, maybe some user doesn't want to expose the fact that they, went to the dotxx TLD

or something. There's tons of TLDs out there. Right? And they may not always be using a large resolver where they can hide in a crowd. Some people run full service revolvers on end stations, some people, you know, run, on a resolver, which has a population of 5 users. So there is a concern. It's not a It's not a major concern, I think, in the community, but some people do have this concern and I think we need to acknowledge that.

TERRY MANDERSON: Terry Manderson, ICANN I think it's fine to acknowledge it. I wouldn't spend time working on it. I think Hans Petter's observation that confidentiality as applied to the systems in routers, all that sort of stuff within the RSS, is by far more important. And that's, I think, that's where the text needs to change.

RAY BELLIS: Ray Bellis, ISC, so the only data that's routinely made public from the root system is the data collection, from DNS-OARC. The source IP addresses on those are all anonymized. So, you're correct that there are people with resolvers who may not be hidden in the crowd, still shouldn't generally be possible to actually identify a single end user. From that kind of data. Yes. We do it. We don't use it.

PAUL HOFFMAN: Paul Hoffman, disagreeing with Ray because this is about incident reporting, not about the yearly diddle thing. So if in fact a root server operator with some confidentiality expectation, and I'm not saying you have them, releases a bunch of queries that might be considered they should have been confidential, but that should have comes from the outside world, not from in here. I think it's worthwhile in this document to specify how much confidentiality is expected to be kept on query data. Sitting at the root server. I hope that the answer is none. I'm not sure that that's gonna pass. But if you don't say it, then and someone goes, well, I saw that such and such happen even though we'd all roll our eyes, this is not so that we don't roll our eyes. This is, as Robert has said a few times, this is to keep the outside world feeling confident that we are watching our own systems and that we will tell them when something has happened. So we need to know what to tell them what has happened, and confidentiality is probably the easiest of those three things. For the outside world to understand about the root server system. And so if we say, and here's what you should understand about our confidentiality. It's very, very low for the queries but it is reasonable for the

people running our routers, whatever. Perfect. But don't assume that the outside world understands confidentiality define it in here.

JEFF OSBORN: Well, that was all rather rousing. We're closing in on the hour. It always feels like a long day when you work on a Sunday. One last thing I wanted to do was I'm hoping that as part of being at IETF, you get to meet people eventually. This is still the 1st day. In fact, it's kind of day 0 or day -1 in a sense. So, I thought during the week, I find it always interesting to pull somebody aside that I've had something interesting to say or something I wanted to know. And IETF is a great place to do that, but you can't do it if you don't know who the people are.

So I would like to ask everybody in here who is a member of either root ops or the RSSAC for their, or actually, this is an RSSAC meeting. Okay. Everybody who is a member of RSSAC means they represent a root server 1 of the 13 and can answer questions this week. Let's all stand up. RSSAC. Okay. So any of us have 12345. This nearly half of of the group. If there was a question about ask an RSO, is there something you wanted to know If we don't know, we can point you to the right people and you have a whole week here. So we're trying to become better known, but what we do is kind of complicated, easily misunderstood and the long conversations that IETF allows for are really good places to learn that stuff. So, with that, thank you all for attending, and Ozan?

OZAN SAHIN: Yes. The welcome reception is starting now, for in person participants. Before you leave the room, if you haven't found your name on the attendance sheet that was circulated in the beginning of this meeting, please just mark you name on the attendance sheet. Thank you.

JEFF OSBORN: Our next, in person meeting is at IETF 120 in Vancouver in July. And with that, we are adjourned.