ICANN69 | Community Days Sessions – RSSAC Work Session 1
Tuesday, October 13, 2020 – 16:00 to 17:00 CEST

OZAN SAHIN: Hello and welcome to the RSSAC Work Session on the Rogue Root Server Operator Work Party. My name is Ozan Sahin and I am the remote participation manager for this session. Please note that this session is being recorded and follows the ICANN Expected Standards of Behavior.

During this session, questions or comments submitted in the chat will only be read aloud if put in the proper form. I will read the questions and comments aloud during the times set by the Chair or moderator of the session. If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor. Please state your name for the record and speak clearly at a reasonable pace. Mute your microphone when you are done speaking. With that, I will hand the floor over to Ken Renard, the work party leader.

KEN RENARD: Thank you, Ozan. Welcome to the meeting on the RSSAC Rogue Operator Work Party. Just curious, Ozan, did you want to do a roll call, list the participants for the record or not?

OZAN SAHIN: For the interest of time, the attendance will be taken from the Zoom room this time.

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

KEN RENARD: Great. Thank you.

So again, welcome and hopefully get some good discussions going here. I wanted to … Here's the agenda in the Zoom session. Quick chat about the summary from last meeting, about a discussion we had in the last meeting and on the e-mail list about where we get our zone data from. Some new examples in the document and revisiting some of the unofficial responses versus rogue responses and a sort of renewed focus back onto an actual rogue RSO. And hopefully we'll get some volunteers to do some writing assignments. There's a few small ones and they should be pretty easy. So, if anybody has anything else to add to the agenda or thoughts, please raise your hand, or speak now, or put something into the chat and we'll welcome any changes or additions to the agenda.

With that, we'll go into a quick summary of our last meeting. We basically went through this document. Ozan, if you wouldn't mind, if you can put a link to the document in the chat. I think most people have it but it would be a good reference.

We talked about the definition of the rogue operator and basically centered around four of the 11 guiding principles that are set out in RSSAC 037, which is the governance model proposal. So, if you take a look at the document, take a look at those guiding principles, so we decided to put some examples in there.

And if, Ozan, you can bring up the document, we can look at … I am looking at a comment here that I made. Starts with "notes from a previous call." It might be down a little bit further towards the rogue RSO scenarios. Yeah, past the background section and a little bit further down. And it should be notes from previous call. Can you expand that comment? Back up a little bit. And there. It's the comment on the right there. Perfect. That's it.

So, these are the ideas that we had from the last call as examples of scenarios to further give the reader a good idea of what we mean by a rogue RSO. So, I went through and put some text into each of these. We can absolutely argue whether there should be more or less said about these, with the one caveat that scenario number four, I really didn't add too much. And there was a comment in the document that I think is very good.

So, scenario number four is failing to conform with the DNS specifications. And my example was improper encoding of a message. You have to include things like misinterpretations of response codes, things like that. But the comment made was specifically towards RCODEs or EDNS0 values that might be incorrect, or if there is any ambiguity that could be misinterpreted there.

So, for … I guess I'll just start going through the scenarios and if we can go back up to Scenario 1. So, this scenario is an RSO, basically deleting a top-level domain—so, taking a few records out of the root zone to effectively make a top-level domain disappear. Now Fred, you might be able to comment on this. This was, I believe, an example scenario from a constituency within ICANN that maybe had this concern.

So, the idea here is that this violates the guiding principle number 2 that says the true source of the root zone is IANA. If you're deleting something, it's no longer the exact IANA zone. This absence should be detectable to validating, which would effectively get an incorrect answer due to NSEC invalidation and hopefully the resolver would then go to a different root server operator.

Any thoughts or comments on this scenario, whether we should add more text, add more detail? Is this an appropriate or representative scenario?

FRED BAKER: Well, you asked me to comment on this. This has been brought up by two parties. Mr. Putin from the Russian Federation gave the possibility of incorrect Internet services and access as a reason for Russia to take control of its borders and operate a separate Internet.

And recently, I was asked, as Chair of the RSSAC, to talk with a Chinese organization about the possibility of a root server doing something terrible. And specifically, they were worried about someone removing .CN from the name service. And in both cases, while of course, if we're delivering the IANA root zone, the IANA root zone contains the Russian Federation and contains China so that's going to be there. That's not something that we have the capability of doing, other than if we operate in a rogue manner. So, that's basically an answer to your question.

KEN RENARD: Okay. Thank you, Fred. So, this is a specific concern that was brought up by someone and I think serves as a good example. Now in giving examples here, the notes from the last meeting give some target scenarios. In fact, the first three scenarios all reference modifying the IANA root zone. So, disappearance of a top-level domain, stripping DNS signatures, or modifying NS and glue records. Thank you, Paul Muchene. Oh, Paul. I'm sorry. I don't see my hands. I will fix that. Paul.

PAUL HOFFMAN: So, I am not clear, Ken, on whether you want us—whether suggestions for adding scenarios are appropriate here. You said rewording the scenarios but are … Is this a reasonable place to be also adding scenarios?

KEN RENARD: Yes, I do believe so. As far as we could come up with tons of scenarios and probably bore the reader to death with them but my—

PAUL HOFFMAN: So, I wasn't suggesting coming up with tons and boring the readers to death. I'm suggesting that, as Paul Muchene had said, maybe Scenario 3 is modifying NS and glue records. I actually don't think that that's appropriate to mix those in a scenario. I would like our scenarios to be very clear on changes. So, my proposal here was going to be to add a second scenario, which is not disappearance of a top-level domain, but in fact changing the NS records.

So, disappearance of a domain, as Fred has said, has been brought up by a number of concerned parties, not just countries. But changing the NS records which is quite different—saying that the zone exists but this party over here is going to be—this different party who's not the right one is going to be doing it. I think that that's a separate scenario. And I think it's a separate scenario from Scenario 3 which is modification of the glue records because modifying NS records is also breaking the DNSSEC signature whereas Scenario 3, modifying glue records, does not.

KEN RENARD: Correct and that's a good point. Overall, I would like … I think that we should have a scenario that represents at least each of the guiding principles—a violation of each of the guiding principles. Whether we have more than one scenario that defines—that talks to a specific guiding principle violation—that's fine.

I am leaving it—sending it out to the group to decide should we have … How many should we have, what are appropriate? And if you think that a particular scenario is worth mentioning here, would you volunteer to write some text for it?

PAUL HOFFMAN: Certainly.

KEN RENARD: That would be great and I will—

PAUL HOFFMAN: I'm sorry. This is Paul Hoffman again saying certainly. I won't say that for other people.

KEN RENARD: That would be rogue. So, if you wanted to contribute another scenario, that would be great.

So again, we can talk about addressing specific issues that are brought up. Specific concerns within the ICANN community is a good target for describing a scenario, as well as just significant events. And I think some of those are referenced here. So, let's see. So, stripping the DNSSEC signatures, I will put that in here and go over that here. So, if an RSO removed signature records from a zone, again, removing any record from the zone or modifying the zone in any way is a violation.

So, at this point, validating resolvers should effectively be denied service from that root server since the signatures will not check and, of course, nonvalidating resolvers would not be affected by the stripping of the DNS signature. Any thoughts? I see Brad's got his hand up. Brad, please.

BRAD VERD: Hey, Ken. The way the document reads right now is a little—I don't know—doesn't flow as well as, I guess, if I'm an outsider viewing this. And since you've pointed out that scenarios one through three are tied to modification of root zone data outside of IANA, is it reasonable to think or could we move these scenarios as examples under each of the guiding principles since you're trying to tie it back to each guiding principle you've called out here? So, back up where you have definition of rogue operator and then you've got Guiding Principle 2, Guiding Principle 6, 7 and 11, under Guiding Principle 2, you would put Scenario 1.5, 2 and 3 right there in the document, just so you're not jumping around to see it.

KEN RENARD: That sounds very reasonable to me. Does anybody else have any opinions on that?

BRAD VERD: So essentially, the document would read the definition of rogue operator, Guiding Principle 2 where you state what it is, and then just below that, "Here are some examples of rogue behavior that is—that goes against Guiding Principle 2," and then the same for 6 type of thing.

KEN RENARD: I like that idea. I will take that on as just the, essentially, reformatting part of that. But yes, thank you.

So, on to Scenario #3. Again, we're talking about violating Principle 2 about the content of the zone. I'll just give you guys a chance to read through that and please let me know if you have any comments or thoughts on that.

Duane, yes.

DUANE WESSELS: Hi, Ken. Thanks. So, one thing that comes to mind as I read these is that they generally talk about modifying the zone or the zone is no longer the IANA zone. But of course, name servers operate on a per-query-and-response basis. So, you can take these actions on individual queries and responses and it's kind of not the same as modifying the zone.

KEN RENARD: Okay. So, maybe more clearly, it would be modifying a response or a response that does not reflect the contents of the zone?

DUANE WESSELS: Well, I guess but then I feel like we're kind of venturing back into the territory of rogue responses which we've kind of pushed out of this document, right? So, I'm not sure what to do about that but I just wanted to point that out, I think.

KEN RENARD: Okay. Paul?

PAUL HOFFMAN: Following on with what Duane just said, one of the things that people were concerned about, and that specifically came up in the earlier discussion, is that a root server operator, for whatever reason—and again, without going into intention—might send out different responses to different queries, such as based on the AS number of the query. So, and this goes back to the question of they haven't modified the zone. If you're from any other AS, you're getting a correct answer. But if you're from a particular AS or a particular set of ASes, you are getting a wrong response. I think in that case, it really is better to focus on responses, but to tie them back to the originator RSO. So, the RSO giving a wrong response is a rogue operation. Thank you.

KEN RENARD: Okay. I think, at least from my perspective, how I'm thinking about this, these are not necessarily just details but they're cases where maybe I modify the zone only for this one response. Okay, that's semantics but one of the other guiding principles is to be neutral and impartial. I interpret that as … Here, "an operator shall respond to queries without bias to the source or content of the query." So, if I give the same query a different response based on who the querier is, that's not being neutral or impartial.

So, I definitely understand what you're saying and I see that. So, focusing on responses, I will take that on versus modifying the zone. I'm writing down my to-do list. Okay. Okay, I will attempt to clarify that, as well as putting these into the definition section. All good ideas.

So, on to Scenario #4 here. Again, I talked about an RSO responding to queries using maybe the wrong RCODE or EDNS code. Thank you, Paul, for bringing that up.

So, kind of looking for a real good example here. I just threw this text in as trying to throw my idea into the document real quick. So, anybody have thoughts on a good example that maybe looks like an RSO trying to do something bad and they could do something bad by maybe changing this RCODE, changing this field, or doing something within the IETF-defined specification for the message formats? Basically, looking for somebody to turn this into a much better example. Paul, please.

PAUL HOFFMAN: So, I'm not sure if we can go with the current wording for your request. There's a question of being conformant with the DNS specifications and then there's a different question of conforming to the DNS specifications but giving responses that would, in fact, have a rogue-ish—a negative effect for some users.

So, conforming with the DNS specifications is then, as you pointed out in this, that it could be just the wrong encoding of a message, but I think really what you might be trying to capture here is not conformance but misbehavior using the specifications other than the root zone data. Is that correct?

KEN RENARD: I believe so. So, what I'm thinking here is we have the RFC that defines specific fields. It could be a semantic misinterpretation of an RCODE or an EDNS option that is maybe outside the specification of the RFC, either a misinterpretation or a flat-out lie. I think of the—

PAUL HOFFMAN: So, I think those two are very different. A flat-out lie that I can think of in this case would be sending a serve fail RCODE to a resolver who somehow you know, if they get a serve fail, they're not going to look beyond your RSO. So, a serve fail is different than an X domain but it might have the same effect for some people. That's still conformant to the DNS specification. It's just a lie.

KEN RENARD: Okay. Again, we're looking to get an example here of violating the idea that the IETF is defining the technical operation of the DNS. So, any of those examples that really hit that point—changing the title of this scenario, changing any of the text … Again, this was my attempt to throw ideas onto a document to start the discussion. Paul and Wes have their hands up. Okay. Wes?

WES HARDAKER: I think Paul is right with … I agree with the things he was saying. And it really comes down to you have to decide is the response that is bad intentional or is it in error? And catching that is hard. It sort of requires detecting different responses to different clients or different responses per TLD. So, you can … And some real quick ones I was thinking of that were not conformant with the IETF specifications as agreed to by the RSOs would be not supporting the EDNS0 max message size parameters by a client and sending them something back too big, right?

And now whether the intent there is to overwhelm their buffer size, given the fact that they told us what it might be, or whether it's just a mistake because some server didn't support EDNS0 properly. Or the other obvious one is not honoring the [dubit], right? They wanted DNSSEC compliant information and some RSO deliberately didn't return it to possibly a particular client or for a particular TLD and stripped it. And that would be, of course … That's definitely rogue if it was done intentionally. If it was done for all TLDs and for all clients, then you kind of wonder, "Okay, well, then there's a bug in their software."

KEN RENARD: Thanks. Yeah. So, the [dubit] would be a great example here. If, I guess, per RFC—if the client sets that a server should/must—I don't know what the verbiage is in there—shall reply with DNSSEC records. So, that would be an example of nonconformance with the specification. In that case, there's nothing to interpret there. It's modifying behavior.

That would be a good example here. Is anyone willing to write up a few sentences using that, using the [dubit] as an example? Wes, is your hand up again?

WES HARDAKER: Yeah, I made the mistake of leaving it up, so therefore, I'm volunteering, I guess.

KEN RENARD: Yeah, I was going to go there.

WES HARDAKER: I can try and write something up. You've got to give me a couple of weeks. I'm already backlogged on other things I need to write up for RSSAC. But yes.

KEN RENARD: Sure, thank you. And feel free to change the title of the scenario as well to best reflect what you say. The idea here is Scenario 4 is looking to give an example or a scenario of where an RSO is going rogue by violating Principle #6 that says the IETF defines the technical operation of the DNS. So, thank you, Wes.

On to Scenario 5. These two, 5 and 6, have really not been thought out thoroughly at all. They are waiting for a volunteer to write a few quick sentences about a scenario. So, for Scenario 5, we're talking about some scenario that we're violating Principle #7, the ethos integrity or working for the good of the Internet.

So, there's a couple comments here in angle brackets of possible ideas that could be used as a scenario example. Disrupting or subverting the work of ICANN is one, using malicious routing tricks, or disparaging other RSOs. So, those are a few ideas to turn into a scenario description. And if somebody would like to take that on or has comments or even other ideas as a good example of violating that principle, please feel free to join in and hopefully write some text. Yes, Paul?

PAUL HOFFMAN: So, I would like to remove the first of those proposals as violating Principle 7. ICANN is a community-based organization and the only way we can change is when the community asks us to change. I don't want any RSO to think that disrupting what we're doing or subverting the work by suggesting changes or saying that ICANN … By the way, I'm sorry. This is Paul Hoffman. I'm ICANN staff but I certainly don't speak for all of ICANN.

But I don't want any RSO to feel that they would be inhibited from criticizing ICANN in the community or trying to change the way we work. I don't want them to worry that they would then be considered rogue. So, I propose just removing that first one. And all of us at ICANN are reasonably good at listening to community suggestions even if it's things we don't like. So, let's just get rid of that one. The other two do actually seem like they would violate Principle 7. Thank you.

KEN RENARD: Okay. And I have no problem with that. Again, these are thoughts. Pick one of these and we can write about it. My initial thoughts of disrupting or subverting the work of ICANN would have been, say, within an RSSAC meeting trying to keep people from getting work done. I see Terry has his hand up.

TERRY MANDERSON: ICANN staff and root server operator. I'm just wondering has any discussion happened around Scenario, or sorry, violating Principle 7 on the actual data that a root server operator sees? At the moment, while we have not fully ubiquitous deployment of QNAME minimization, if that would ever happen, there is value in the data. Even with QNAME minimization, there is still value in the data. Some thoughts, perhaps there, might be interesting. Thank you.

KEN RENARD: So, are you proposing a topic for Scenario 5 that could be violating privacy—selling the data that an RSO might see, something like that?

TERRY MANDERSON: Something like that.

KEN RENARD: Okay, I'm going to try and capture that as well here. Misuse of data collected.

Again, so these are all ideas. Pick one, and go with it, and write up some scenario. Again, it's only a couple sentences. So, I'd really like to have a volunteer to pick one of these and go with that. Terry, you should put your hand down or I will interpret that as volunteering. Darn. Okay. Paul Muchene?

PAUL MUCHENE: I think I can volunteer for the first proposal, using malicious router tricks. I can write a sentence on that.

KEN RENARD: Thank you very much. That would be very helpful. And go ahead and add that to the document when you can, just knowing that section may be moved up into the previous section at some point.

Okay, on to Scenario #6. Again, the same idea here. We need somebody to write a scenario for—that would demonstrate violating Principle #11, the neutral and impartiality. My thought was do not respond to any requests from a certain country, as Paul adds, or a specific ASN. I also really like the idea of the second example, implementing DNS policies of a specific government. To me, that's a really good example of neutrality, impartiality, and especially within the less technical parts of the ICANN, the more political side.

So again, those are ideas. Looking for other potential ideas to describe a scenario that violates Principle #11. Anybody have any thoughts or anybody willing to write a few sentences about a scenario? Looking for hands. I wish there was a roulette wheel that we could spin and pick somebody. Okay. We'll come back to that later on when we try to discuss writing assignments. Wes?

WES HARDAKER: Yeah, I think that's going to be a hard one because it's very similar to the other one I already offered to create some text for. I guess it's similar to most of them, right? Neutral and impartial is hard to figure out exactly how that would … Anything that you come up with here will not also be placed in one of the other scenarios.

KEN RENARD: True. If you're … What I can think of as far as implementing policies of a specific government would also be modifying responses or modifying the zone.

WES HARDAKER: Right and you can maybe do something like rate limiting or something. But Paul Hoffman's going to have a better idea. He has a hand.

PAUL HOFFMAN: No, I don't. I was going to agree with you, Wes, that there is a large amount of overlap. And I think that's fine. Given what Brad had suggested earlier—that these are not going to be standalone sections, that are going to be lists underneath the guiding principles, I think it's fine for the list under Principle 11 to simply refer back to other examples that have already been given. Since it's the last principle, I think that that would be easy to read is to say, "For example, like we set up above here and like we set up above here." I don't think that they have to stand on their own.

KEN RENARD: Thank you. Yeah, I think the violating Principle #11 had a lot of interpretations and it speaks maybe more to the intent than the action. Okay.

One of the other things that we talked about in the last section, again, for readability of the document and really focusing on a real operator going rogue—a real operator doing something bad versus intermediate responders and things like that … So, I moved the section on damage that can be done by a rogue operator up before any discussion of what we'll call for now unofficial responses.

So, this is all text, in the next two sections … Damage that can be done by a rogue operator and detecting and mitigating rogue RSO behavior are two sections that effectively moved up in--up to this part of the document and tried to expand on. If you can, take a glance through that. Looking for any comments or thoughts on that.

The detecting and mitigating a rogue RSO behavior section, that is just rough draft text, ideas thrown in the document for the sake of discussion here. And I'll give everyone a few minutes or a few seconds to read through that.

Paul, yes?

PAUL HOFFMAN: I disagree with the second sentence in the first, in detection, that glue modification is not directly detectable by validation of the root zone but are sometimes detectable when resolving the delegated zone. It is detectable if you're also holding a copy of the root zone and you can compare. That's exactly how we were intending to do it with the RSSAC 047 test.

So, I mean, it depends on who you're saying wants to be detecting it. If you're saying a resolver should be responsible for detecting, I think that that's misplaced. If it's some outside agency who is looking for—who is charged with detecting, detecting modification of glue records is trivial.

KEN RENARD: Good point. Yeah. That would be a specific measurement which is certainly within the scope of trying to measure and detect a rogue behavior. I just added the comment there, extra word, glue record modifications are generally not detectable. For a typical user or for typical operation, they would not detect it at that point. But yeah.

PAUL HOFFMAN: But so, but Ken, I want to press a little bit harder on this. Are we expecting them to be the ones detecting? I don't think that a resolver operator is the right one to be determining rogue-ness partially because it would be easy for them to spoof their answers in order to, for example, hurt an RSO. I think that you need to decide who is meant to be doing the detection and mitigation before we say what their capabilities are.

KEN RENARD: Okay. So, all right. So, the glue record modifications are not directly detectable by DNSSEC validation of the root zone but are detectable by direct comparison or by resolving at the delegated zone. Does that more accurately reflect?

PAUL HOFFMAN: I see other hands who might want to weigh in on this.

KEN RENARD: Okay.

PAUL HOFFMAN: Yes, it's true. But again, I don't know whether it's relevant.

KEN RENARD: Okay. Brad?

BRAD VERD: I think Wes was up first. I'll defer and I'll come after Wes.

KEN RENARD: Okay. Wes?

WES HARDAKER: All right, thanks. Paul brings up a good point that that sentence has a problem because glue records are detectable, assuming a couple of things. Adding to his, you need … It's detectable if you have an authentic copy of the root zone, not just a copy of the root zone. It's got to be one from a—that you've authenticated through some other mechanism. It might be you've pulled it over HTTPS from IANA, or it might be that you've gotten an AXFR that was properly authenticated, or that you have used what we might hopefully put in the root zone someday with doing checks on proposal, for example, that was then verified with DNSSEC.

But the important part of that is detection by whom is a good point too, right? I think part of this point is can a client be easily deceived? That's sort of the point of the document. If an RSO goes rogue, can they deceive clients? That's a different question, than can they deceive auditors? And I would say that both of those two points are equally valid to consider. But when we're thinking of problems to put in here, we might, we could separate them out and say, "Well, an auditor might be able to more easily verify that this is not being done."1 It would be much harder for a client to notice.

KEN RENARD: I think that's a good point. And if we can find somebody to expand on this section using these thoughts. Duane volunteers. Thank you very much, Duane. The only thing I would do, Duane, if you can, just maybe in your comments, reference these discussions so that we can, so they can be captured. Brad and then Muchene.

BRAD VERD: I was just going to add pretty much what Wes said but I think … I feel like some of this is where we, again, as a group, try to boil the ocean and put all this stuff in one place. I don't feel this document was ever intended to tell a resolver that they should be detecting rogue behavior. I could easily, in my head, think through a scenario where somebody is complaining about a root server operator going rogue. And then, using Wes's terminology, some auditor would come in and check that, probably in pretty quick fashion since going rogue is means for dismissal. And I think that auditor would probably be the SAPF. But I feel that we … Let's not lose sight of the document of defining what is rogue and not necessarily defining who's going to be doing it, if that makes sense, or how it's going to be done. Or if so, then we need to change the scope of it.

KEN RENARD: To follow up, I agree with you completely, Brad. I was not trying to say that we should alter the scope or the—

BRAD VERD: No, no, no. I didn't think you were. I just feel like as a group, this is kind of what we do in these different scenarios as we talk through them and we just continue to add to this document. And again, the intent of this document, I believe, was what can a current root operator do that is rogue behavior? And one of the reasons we wanted that is because we called out rogue in 37 and now there's an SLA/LOI document out there that basically defines by going rogue and not remediating it, in a certain amount of time, you can be removed.

So, I think this is important that we just don't lose sight of what this document was supposed to cover. And if there's more stuff we need to cover, maybe we put it off to a different document. But I don't want to lose sight of that.

KEN RENARD: Okay. I think a small section on detection, maybe we could even strike the mitigating. From the reader's perspective, I think a lot of the mitigation or detection is going to be centered around, "Use DNSSEC." And if that's a subtle message of this document, I think that's fine. Muchene?

PAUL MUCHENE: Hi. I just wanted to add to that. As far as DNSSEC signed zones are concerned, the dual record modification could send the resolver onto a different path—could send it to a different name server. But DNSSEC concerns itself not with where the response came from mostly, but whether the response is authentic or not. And so, in this case, as far as validating resolver's case is concerned, a true record modification unless it actually modifies the data as well, would not be a problem. But if it modifies the data—if there is a rogue party somewhere that modifies the data, then it will always be detected.

So, either it modifies the data, or it does not respond properly, or it's unavailable, that particular address to address record … In any case, the resolver will always see a failure—the validating resolver—especially coming off the root zone. So, it's not that they're not detectable. For somebody who is, I think, running an authentic resolver service or authentic resolver operator, they will notice a problem, a root response does not validate or a response chaining off the root does not validate. And so, they will question why this is happening because that TLD will not work anymore. So, it's not that it will not be detectable. It will just be problematic.

KEN RENARD: Right. Thank you for that point. And in, I think, many cases, a resolver, if it gets a detectable error, an invalid signature, that resolver could choose to use a different RSO. So, the net effect on the user is really just a delay. So, thank you for that.

So, Duane, thank you very much for volunteering. There is some good stuff we can do with this, not getting too far out of scope. But I think showing that these errors are detectable, giving an idea of what might be done to control this, may give the reader an easier feeling about how things go rogue.

So, we got about ten minutes left. I wanted to comment here and maybe just set this for homework or thoughts for next time. The discussion of unofficial responses. We used that terminology early on in the work party. As we've refined the document, we've refined what it means. And I wanted to share my thought process on this discussion of why it's important, and what it really means and maybe how we change it to better reflect what we mean.

So, way back when, if you do look at the rest of this document, there's a lot of stuff about different types of scenarios where somebody else is responding on behalf of the RSO and this attack versus that attack. I'd like to basically boil all that down to maybe just this one paragraph or something like this. So, I throw out this paragraph that's right now titled "Unofficial Responses" as a complete replacement for the entire discussion of this unofficial response topic.

Another aspect of this is maybe some confusion that could come up with the term "unofficial." If we were to use a term that was even maybe more accurate, we could say a "non-RSO" response. So, this is a … So, we define here … I'm going to read the, I guess it's the third sentence of the unofficial responses section. "An unofficial response is considered anything other than an authentic response from a system authorized to operate a root server IP address in the root zone." So, think of this as an IP packet that's not the one that came from the RSO or a TCP session not from the RSO.

So, a couple things here. What do people think about changing the term "unofficial response" to just "non-RSO response?" Are there any other suggestions that—for terminology here? really my idea behind this is to make sure that any detection or identification of a rogue operator is accurate. We want to avoid false positives and false negatives of a detection of a rogue operator. I'll give one example here. If an observer is trying to detect a rogue operator or doing something and thinks that they detect a rogue operator, but the response that they're actually getting is an incorrect response from an attacker, that is something that should not be used to determine the rogue-ness of an operator.

Another example might be an 8806, the local root. The local root could be providing a correct answer to the observer but maybe the RSO is actually rogue and would have responded with an incorrect or modified response. In that case, the observer sees what looks like a non-rogue RSO, but in fact, it is rogue. So, the context here is how we measure or identify a rogue operator, technically. So, in this time, I'd like to hear some thoughts on that and pardon me if I'm out of order here but I'll go with Paul first.

PAUL HOFFMAN: So, I like the idea of calling this non-RSO. I think that that's much more accurate. I strongly object to mischaracterizing RFC 8806 responses here. RFC 8806 is only about resolvers. It has nothing to do with authoritative servers. So, if a stub resolver has sent a query to an RFC 8806 resolver, it is not talking. It never expected to talk to a root server or to get a response back. So, please don't bring that up. It's confusing. It goes against the RFC. But there are plenty of examples, as you have said, of non-RSO responses which we do care about. Thank you.

KEN RENARD: Okay. Wes?

WES HARDAKER: Thanks. I wasn't going to talk about 8806 but now I have to. So, Paul's right that people querying the 8806 resolver aren't querying the root so they wouldn't expect to get something. That being said, the 8806 resolver is actually querying itself or a parallel authoritative server, in some cases, where the resolver actually isn't able to serve itself. So, there's that mirror instance that's used by at least the earlier versions of unbound, for example.

But I think what I was really going to say is I like this whole approach in the first place and that we're only talking about it. I will say there's a few times that it's actually hard to determine whether or not you're talking to the root server system per RSOs directly. And a classic example of that is a paywall type environment, especially in hotels and stuff, where they are deliberately doing man-in-the-middle DNS answers. And you can detect that for very—TTO watching and things like that that show that even though you're sending a query to the root, it's being intercepted and actually handled locally.

KEN RENARD: Right. So, that is a good example of a non-RSO response. In the context of mentioning 8806 here, I kind of go back to my example. If I'm,1 as an observer, am querying my enterprise resolver and it's using 8806, a query never goes to an RSO. It uses the data that it obtained via transfer. Could that …? Since it's not an answer from an RSO, could that potentially hide the fact that an RSO is rogue but since you're not querying the RSO, you don't see that?

WES HARDAKER: Well, if you're not … So, I think Paul probably has his hand up and would answer. But I would say if you're not querying the RSO, then you don't expect a response from the RSO. That's what Paul was saying. What I was saying, the resolver may query what it might think is the RSO. Paul ought to probably answer that.

KEN RENARD: Right. In that case, I think the observer may think that they are querying the RSO versus that expectation that it's not querying. Paul, please.

PAUL HOFFMAN: So, we cannot say that a resolver is expected to give the same answer as an RSO simply because, for example, I think it would be considered rogue if a resolver sent a query for the NSEC for .COM to an RSO and got back a TTL of 60. That would be considered rogue. That's not what the data in the root zone is. However, that's perfectly reasonable for a stub resolver to ask its local resolver what is the NSEC for .COM and get back a TTL 60 because that's all that is left in that resolver's cash.

So, Ken, I really object to the idea that we can make resolvers, whether they're doing 8086 or not, be part of the system where one is expecting responses from the root server operators. I think if you want a response from a root server operator in order to detect rogue-ness, you got to go to them directly.

KEN RENARD: Okay. So, if we were to just … An observer that's just sitting on their laptop with maybe no understanding of their enterprise recursive resolver, whether or not it's doing 8806, they could observe something, thinking that they've queried the root zone, but in fact, they have not?

WES HARDAKER: Ken, let me interrupt because I've got a hard stop that I've got to leave immediately. You need to think of it in terms of who is sending the query. Is it a client and are they trying to reach an RSO? Full stop. That is the definition of rogue. Do they get a response from something—an official RSO or an unofficial RSO? And just think of who sent the query. The observer … I like the idea of thinking about observers but that actually makes it, I think, a little bit harder to think about. With that, I apologize. I've got to leave.

KEN RENARD: Okay. Thanks everyone. Sorry. Yeah, we are getting close to time. I think, Paul and Wes, I am starting to understand what you're saying. It's a different way of thinking. Let's bring that up for the next call. I think, Paul, maybe I'll continue, if it's all right with you, to discuss this with you and try to put something in here that more accurately reflects that.

So again, sorry about the running pretty late here. Ozan, if you wouldn't mind putting up the agenda that shows our next meeting time. I will try to send out a summary of today's call. And the next work party meeting is Tuesday, the 17th of November. That is actually during the IETF week. Anybody have a conflict with that? Should we move that to week before or week after?

UNIDENTIFIED MALE: We can't know the conflict until the DNS-related working groups have their schedules and that won't be for a few weeks.

KEN RENARD: Should we push it in the anticipation, however unlikely or likely, might as well just to avoid a conflict?

STEVE SHENG: I would recommend we push it by one week. I think it's probably good to respect IETF's schedule that week.

KEN RENARD: Okay. Yeah, so let's do that. Let's make that, I guess, the 24th of November and is that …? That's the Tuesday before Thanksgiving in the U.S. And with that, thank you for joining us today. Thanks for the discussion. Please feel free to comment in the document, to comment on the mail list, and anybody that's willing to write some text for any of these sections, please go ahead in the document and we can discuss next time. Any closing comments, thoughts from anyone? All right. Thank you all and I guess back officially to you, Ozan.

OZAN SAHIN: Thank you, Ken. Could you please stop the recording, [Moses]?

[END OF TRANSCRIPTION]