DANIEL GLUCK: As Geoff said, welcome to the first RZERC meeting of the year, #48 in the, I guess, series. It's Monday, the 27th of January at 1900 UTC. This call is being recorded and is subject to ICANN's Expected Standards of Behavior and Community Anti-Harassment Policy. And with that, Geoff, please get us started.

GEOFF HUSTON: Right. I don't think it's worth going through a roll call. It just takes time. I not that we have apologies from Liman and Wes Hardaker. The only other person I can't see here at the moment is Carlos from the ASO. And everyone else is here. So for that we can assemble in the minutes. The list of people is here, Dan. I'll leave that up to you. And that means that we can go straight on to Willem with the recommendation response from the study in RZERC002.

WILLEM TOOROP: Yes. Let's see. Can you see my slides?

GEOFF HUSTON: Yes.

WILLEM TOOROP: Excellent. The Root Zone Evolution Review Committee discussed the DNSSEC signed root-servers.net zone in response to the Root Zone Maintainer representative proposal to signed root-servers.net. And the RZERC decided it needed further study, and RZERC002 report states

RZERC's position and two recommendations; a consortium of NLnet Labs and SIDN Labs, which is part of the NLnet Labs, have worked on both those recommendations as commissioned work from ICANN. And this has resulted in two reports.

So the first recommendation of RZERC002 was to conduct the further studies called for in Recommendation 2 of RSSAC028. We've done that, we've presented about this too. The report is called the RSSAC028 Implementation Study Report, and you can find it on the ICANN website.

So this is a more detailed view of the recommendation. It is, "Revisit the options and consequences of signed root server data; the behavior of authoritative nameservers with respect to response sizes and truncated responses; and also, understand and document the behavior of recursive nameservers with respect to validating signed priming responses.

Yeah, so RSSAC028, the original report suggests a few different naming schemes for the root servers to evaluate. In the report, it is evaluated with respect to response sizes, et cetera, what I just said. And this slide and the next slide is an overview as a reminder of those naming schemes. So 5.1 is the current scheme, 5.2 the current, but DNSSEC signed basically, but the Root Zone Maintainer representative suggested, which was the cause of the RZERC002 reports.

5.3 is an interesting one because the names, as well as the addresses of the root servers are authoritative within the root zone. 5.4 is a shared delegated top-level domain for all the root servers. Just taking out .NET,

so to say 5.5 is a scheme where each root server operator has its own TLD, it puts the addresses for the root servers. And the name authoritatively.

5.6 is also a scheme in which everything is authoritative, but then everything is with a single shared label for all the operators. So the RSI's, the root server identifiers, there's only a single root server identifier. The numbers are the sections within the RSSAC028 report, and I'll use those numbers to refer to these naming schemes.

So this was quite a comprehensive study, and there's a large report. So I thought for this presentation, I'll skip immediately to the conclusions. One of the observations was that since the initial report, DNS Flag Day 2020 has happened and it has had a big impact. So not only do recursive resolvers increasingly send requests with EDNS UDP max payload size of 1232, but also authoritative nameserver software returns smaller responses by default.

So most root nameservers currently return responses smaller than 1232, but not all of them, regardless of the EDNS UDP payload size in the request. That is to say, if the requested payload size is smaller than that, then the authoritative will return a smaller response, but not if it's larger than 1232. Fragmentation is perhaps not so much of an issue anymore.

Most root name servers can also return all names for some or all addresses, sometimes with signatures, sometimes without, without the truncated flag set, the TC flag, for all naming schemes within the flag day limit of 1232 bytes.

Some root name servers will not do that for the schemes where the root server names and addresses are authoritatively within the root zone. So I will come back to that later. But TCP is for most schemes and for most resolvers not so much an issue anymore, it all just fits by far in the most cases, except with PowerDNS, which does not follow up when the TC bit is set in the primary response. But I'll come back to that later.

So signatures for root name server addresses. Signatures in priming responses are included only when the addresses reside authoritatively within the root zone, but only by some root name servers. So these are the schemes where the addresses reside authoritatively in the root zone itself. And this is interesting because it did not affect the resolvers at all. They don't care if the signatures are there or not, if those addresses are returned in the additional section.

And actually, when you think about it, it makes sense because in the priming response, there's no way to see if those addresses are authoritatively within the root zone or not. If you compare 5.3 and 5.4, the response for the addresses in the additional section for both those schemes can be the same. I hope this is clear. If not, then please give me a signal later. But therefore, those schemes that put everything authoritatively within the rood zone provide no added security whatsoever. That is basically the conclusion.

All naming schemes had issues with other resolvers that we tested, which were older and newer versions of the most prominent open-source resolver software, including the current naming scheme. New IP addresses are not picked up by PowerDNS if the names remain the same of the root NS set.

5.4 had killed the Knot resolver after priming, which is not that good. 5.3 and 5.6 caused some TC bit responses which are problematic for PowerDNS. The one that did not have issues at all, surprisingly, is the one where each root server operator has its own top-level domain to put root server names in the addresses.

GEOFF HUSTON: Can I just interrupt you there very quickly, Willem?

WILLEM TOOROP: Yes.

GEOFF HUSTON: There's a question from Peter Koch there; is this an implementation or a configuration issue?

WILLEM TOOROP: The issue with the PowerDNS or with all the others? With the Knot resolver, it's an implementation issue. Yes, we've tested several different configurations with different hint files where the priming response had different names than in the hint files or the same names, et cetera. For 5.4, that made no difference, for the naming scheme 5.4. And also, yeah.

GEOFF HUSTON: Okay, sorry to break the cadence. Peter's on a train and can't do it himself.

WILLEM TOOROP: Yeah, no problem. There's far more detail. There's an extensive report on this study, right? It's a 60-page report or so. We go into all the details of all those resolvers, what worked and what did not work.

So this is the conclusion of the RSSAC028 Implementation study reports. And then I'll move to Recommendation 2 of RZERC002 reports, which is, "RZERC recommends that ICANN org further explore the cost benefit tradeoffs and risks of signed root zone name server data. Do the risks of redirected query traffic outweigh the risks of increased operational complexity?" The same consortium that did the first study did this one too, which is also published as a report in the OCTO-commissioned works archive.

We reformulated that a little bit because we realized we cannot actually say anything about increased operational complexity. So we reformulated it as, "Does signed root zone name server data reduce the risks of redirected query traffic? And if so, to what extent?" With some boundary conditions, "Only considering the recursive resolvers to root name servers," so not forwarders or those kind of recursive resolvers, and "only with currently deployed resolver software." Basically, we looked at the 2023 "Day in the Life of the Internet" data that is provided by the DNS org to come to a number of what the reduced risk would be.

I think it's good to come to the conclusion immediately and then later have the slides that explain how we come to these conclusions. Already from the previous report, we know that, "Signed address RRsets in the priming response do not help," because an attacker can just remove all

the signatures without consequence for the resolver. It's indistinguishable from a naming scheme where those addresses are not authoritatively within the root zone.

“Only explicitly queried for root nameserver addresses benefit this,” if resolvers request the authoritative addresses for those root server addresses. And the method will be -- we'll go over the method later, but the overall risk would be reduced with 1.2%, with which we mean, “If P is the chance of a fallen victim to query redirection now, then signed root servers would make the chance: P * (1 - 0.012) larger, or the chance would be than this number, this slightly larger than P number.

But it's not equally distributed. It really has everything to do with to which TTL resolvers cap. And there are several bands that can be distinguished. First of all, if you look at this table, the first row are the resolvers that do not do DNSSEC at all. So the reduced risks would be zero there, obviously. And then there is 0.1% of resolvers in the bottom row of this table that we validate all root server addresses. After priming, they just query all the addresses. This is, by the way, typically the behavior of the Knot resolver. And so this proportion of resolvers as seen in the digital data would be protected 100%.

For the others, if they adhere to the TTL that are returned in the priming response and also when querying directly for root servers.net addresses, then the reduced risk would be 5.9%. But most resolvers do not do that. The 1% is a class of resolvers that I don't know. I don't know which software this is, but it caps TTL to 10 days. The one above there is typically bind, and also the suggested value in RC8767 to cap TTLs to seven days. And the one above that is everything else. So just

for the record, unbound, Knot resolver and PowerDNS are kept to one day by default.

Another conclusion that is perhaps more an opinion of the consortium is that, “Since only explicitly queried for root server addresses benefit from signed root server data, and only 0.1% of resolvers query all root server addresses, Knot resolver, some unbounds and it is configured to do that, therefore, there's limited risk on negative impact if we would be signing now, the root servers.net. And it could be an opportunity to sign now and establish a strong case for resolvers to revalidate root server addresses in the future.”

And we believe that DNSSEC is actually valuable to use in the DNS resolution process. And that it's a pity that in the current design, it does not do so in many cases. Even more so than in other applications, we think that yes.

So, what is the risk exactly? Root name server addresses are currently not signed in priming responses. There's no DNSSEC verifiable statement in the DNS that those addresses belong to the root server identifiers. So they can be potentially altered undetected by an attacker, which can redirect the root zone to a rogue root server. The attacker sees all traffic, which is sort of a privacy issue, and the attacker can modify all unsigned data, which includes all the referrals or all the glue in the referrals, so to say. And because of that, the attacker can also change the referral responses and potentially hijack the whole domain name tree.

Signed root server addresses could in theory change that, obviously. So one thing, the first statement says that root name server addresses are currently not signed, but this is actually not true. There is ZONEMD records now in the root zone, which is a signed statement of all the data in the root zone, including the root name server addresses, and also all glue for all the delegations that are in the root zone. So it's resolvers that have the root local and verify that ZONEMD are protected very well.

So how is this possible? This has everything to do with how data is ranked, the trustworthiness of data in the DNS. So there's two kinds of data. There's infrastructure data that's used by resolvers, and there's the authoritative data that is returned by resolvers. Sometimes the infrastructure data is-- Yeah, we know about that. It's the clue. It's the addresses in the additional section from also in the primary response. And only if resolvers query for those addresses explicitly, they get the authoritative version. Sometimes they can learn it also because it comes in the authority section of responses.

So this is the same list as that of RFC 2181, but I gave the list items different values. All the A values means that all that data is authoritative and contains DNSSEC signatures that can be validated. B is all the non-authoritative infrastructure data. And C, which was not in the original 2181 list, is the data that resolvers have when they start. So it's the built-in root server names and addresses. And CCC is the root server addresses and names from root hints file.

So what does it mean, initial priming? It means that the names of the root server addresses are received authoritatively. They can be signed.

They are currently also signed. So they get A status, but the additional addresses not. Both have a TTL of six days. Repeated priming as described in 8109 doesn't change that. The names are authoritatively returned, but the addresses are not.

So what is in this slide is presented, but the risk is that query traffic can be redirected. So for an on-path attacker, it means it has a 100% chance to redirect query traffic to the root server if it can be on the path of one or more of the IP addresses of the root servers. For off-path attackers, it's a bit more complicated. There's a formula in RFC 5452. If we choose some numbers that make the attack more likely to succeed, then high checking priming response with a constant team of 40 megabits per second, I believe, or four. I didn't make a speaking note. Maybe it's four, maybe it's, I don't know. But 39% of success within one year and 64 within two years.

So if resolvers do follow up address queries, that makes all the difference because then it receives the addresses authoritatively not with a signature, but with a higher TTL because root service.net has a TTL of six weeks. And yeah, that makes the numbers. For on-path attack, it doesn't change the things that much. But for off-path attackers, it reduces the chance on success significantly. Even if those addresses are not the NX side. Oh, I see now in my speaker note the query rate of the attacker for the off-path attack. And we took 45 megabits per second, constant stream of fake responses.

So some resolvers do follow up address queries and some do not. PowerDNS never queries for the root server addresses. Knot reserve always queries for all root server addresses. Unbound and bind query

for missing root server addresses or bind only if there are none in the response. So, Knot resolver is very well protected, I guess, of would be, have a 100% benefit of reduction, so to say, of the risk for query redirection if root service .net would be signed. And for the others, it depends on the follow-up queries after the primary response.

But zero additional addresses in the additional section is not an option because there's one naming scheme, 5.3, the one where everything is authoritatively within the root zone. And some root servers will then return the priming response without any address at all. And that's a hard fail for PowerDNS Recursor. It will then stop resolving altogether. This is one of the schemes that is in the RSSAC028 Implementation Study reports.

So there is an internet draft. This is an old version of it to do delegation revalidation. In this draft, is also mentioned, revalidation of root nameserver addresses. There's also an option, configure option for inbound, hardened refill path. If you set it to yes, then it will do delegation, revalidation, and will also revalidate all the addresses that are returned with, in the additional section of authoritative NS responses.

So this is unfortunately not the complete story. This would be true if all resolvers would follow the TTLs that are returned in responses, but they do not resolve as cap TTL. They lower it or they have a maximum TTL configured. And for this to get insight in how many resolvers or what the amount of resolvers are that capture a certain TTL, we scheduled a RIPE Atlas resolver measurement using all the probes querying for a name with a six weeks TTL. And look at the TTL in the query response.

And we associate then the resolver as seen at the authoritative with the response seen from the RIPE Atlas probe.

So, with this measurement, we could identify 16,000 IPv4 addresses and 11,000 IPv6 addresses for resolvers. So that on the left is a distribution of those TTL caps that we saw. And in the middle is a formula to then weight the extra protection that you would have from the TTL from authoritatively obtained good name server addresses, the ones with the six-week TTL. So obviously for unbound nodes PowerDNS that have a cap of one day, there will be no benefit at all, which is so it has a waiting factor of one. The risk is the same, so to say, with the authoritatively required addresses as addresses from the additional section.

So bind that caps at seven days has a slight reduced risk. The resolver class that I don't know about that caps at 10 days have slightly more increased risk and the 8.4% of resolvers that do not cap TTLs benefit most. So, we will use these weighting values on later slides for the reduced risk of query reduction.

So here are some interim observations. Priming response will not and cannot be protected. Resolvers that revalidate root server addresses are fully protected (Knot resolver and Unbound with harden-referral-path). Some resolvers do query for addresses if they are missing from the primary response. Unbound will query for the missing addresses, bind will query for them if there are no addresses in the primary response, but that will kill PowerDNS. So that's also not a general solution, so to say. We can observe both at the root of those resolvers.

So this is the formula how we came up with the metric of risk reduction. So the proportion of usable authoritatively obtained addresses for the root servers in proportion to the number of usable non-authoritative addresses built in from hints and from priming responses. That is basically the reduced risk. If you have the proportion of authoritatively obtained service to the non-authoritatively addressed. But the protective effect of the authoritative RR sets will last the TTL of the authoritative addresses in proportion to the non-authoritative addresses longer, and their rate of appearance in the DITL captures will also be the same amount, fraction less frequent in proportion to the cache occurrence of resolvers on the internet.

The effect of the usable authoritative addresses and this of the whole risk reduction quotient must be weighted accordingly. Maybe I can say it again.

GEOFF HUSTON: No, I'm conscious of the time. I wonder if we could sort of get closer to around slide 37, your conclusions and recommendations at this point.

WILLEM TOOROP: Yes, yes.

GEOFF HUSTON: Thank you.

WILLEM TOOROP: Yes, here's the bending again that I showed earlier as well. Do the risks of redirected query traffic actually the risks of increased operational complexity? With only signing, the risks can be overall reduced by 1.2% as observed with how the internet responded in 2023 as seen in the DITL data. If that is considered not substantially enough, then no. But the risk is real and with severe consequences, query redirection with query redirection, you can hijack the entire DNS.

Only 0.1% of resolvers query for all root server addresses, not an unbound with hardened-referral-path. Therefore, there's limited risk or negative impact if it would be signed now, and perhaps opportunity to sign now and establish a strong case for resolvers to revalidate wood-silver addresses in the future.

Yeah, there are a few more slides with suggestions for the different kinds of operators. For example, resolver operators should do DNSSEC validation. If they can, they should do ZONEMD, which will protect them 100% from query redirection, not only from the root, but also from the TLDs. Cookies is a very good mechanism to protection against off-part attacks because it provides 64 more random bits, even if not implemented correctly.

Another suggestion is to do revalidation of those root server addresses. If resolvers would do it now with the current naming scheme, then traffic to the root will increase with 30%. With the naming scheme that had the fewest issues in the RSSAC028 Implementation study, there will be only traffic increase of 1.7%. So, root server operators can also make the situation a little bit better with respect to query way direction the risk by providing DNS cookies. Some of them already do, but the cookie

secrets are not synchronized among all the Anycast instances. It would be nice if they would do that.

RPKI, so DNSSEC does not protect against route hijacking. So with RPKI, even if the root server addresses would be protected by signatures, a hijacker of the address could then still see, well, it's a privacy risk, see all the traffic to that root server, that hijacked root server. And this is a list of all the letters at the time of the previous presentation that signed their IP addresses with RPKI. And that's the whole deck.

GEOFF HUSTON: Thanks, Will. That's very comprehensive. I'm sort of anxious of the time, and I want to leave enough time for Roy and the KSK. So I suppose I have one big question, and I'm not even sure we can answer it now. Maybe it's for the next meeting. I understand this study was part of the RZERC recommendation and your initial slides did Recommendation 1 and Recommendation 2 that the study be done. Who should respond to this study, do you think, and any other members of RZERC? What's the view on how a response to this study should be formulated? Is it RZERC's problem or somebody else's problem?

WILLEM TOOROP: Well, yeah, I don't know. Well, I see that the Root Zone Maintainer representative is present as well.

DUANE WESSELS: Well, I just wanted to put on a third option, which is maybe it doesn't have to be anybody's problem. I mean, RZERC asked for the study and the study was delivered and we've received it. And so, maybe that's it.

GEOFF HUSTON: Okay. I think there are some interesting points there, and maybe we show up for the next meeting. In particular, the note, I think it's on slide 37 or so, that said if we're going to sign it, now is a good time to sign it. And if that's the case, that's more than just, well, thanks for the report. That's brilliant. It really is, well, maybe there's a case to sign it and that could be considered.

WILLEM TOOROP: Yes, actually, so I think that signed infrastructure data has value for resolvers more so than for example, signed addresses of websites that will be authenticated by TLS anyway. So yeah, I think it's really a pity that currently delegations or referral responses are not signed. Yeah, different resolver-- Yeah, so currently, resolvers don't do anything with this authoritative data, but if we would sign it now, it would be of benefit for those resolvers to start doing it, to start using that data.

GEOFF HUSTON: Okay. I think that's probably enough at the moment to leave this and then can we revisit this study at the next meeting of RZERC with particular attention on whether there is any follow-up actions for RZERC having asked for this study and whether the study itself is self-

explanatory enough or whether there are any follow-up actions. Is that okay?

WILLEM TOOROP: Yes.

GEOFF HUSTON: Thank you. With that, hopefully I've left enough time for you, Roy. The KSK Monitoring Update is the other thing on our agenda. And so, Roy, over to you.

ROY ARENDS: All right. Can you hear me?

GEOFF HUSTON: Yeah.

ROY ARENDS: That's brilliant. I've turned the camera off for your benefit, if you don't mind. So I've been asked by my colleague Kim Davies from IANA to discuss with you what we've decided on monitoring. And nothing is 100% set in stone to be fair, but let me just give you an introduction of what we're doing.

So in the past, 2017, 2018, any monitoring we did then, to be honest, has been really conclusive. We got some hints, some allegations, some assumptions and a lot of bias. And it turns out that 8145, that's the RFC

I'm referring to, and that's the bulk of the monitoring that we did together with all the root server operators, wasn't really helpful. And I could explain what happened then in 2017. And you can find all this in OCTO 12. That's a report that we've written after the fact.

What we saw was quite inconclusive in such a way that we've asked our then CTO, David Conrad, to talk to IANA and maybe postpone the role for a year. And it was discussed with various partners and that was a good idea at the time. And it gave us time to actually research all the data. And it turns out that it didn't really got us any wiser. And let me tell you why. RFC 8145 had basically four issues, if I remember correctly. RFC 8145 data, you can only see that passively. So only for instance, at a root server traffic level, you can't really query resolvers and see what they return using RFC 8145.

Another issue is that whatever you saw doesn't represent the actual uses behind the resolver. So you could see a resolver indicating an 8145 signal that is basically, I don't know, 10 clients behind it or could have a thousand clients behind this. There is no indication in that signal. Another issue is that 8145 queries, they would propagate through resolvers. So if you have a stub resolver doing all the signaling towards a resolver that we see talking to the root servers, then you get basically different signals and the resolver might not even be actual validating, if that makes any sense.

And the last problem is not that we have evidence for this, but we do know that that data can be spoofed. So if one wants, it can basically influence all the signaling because nothing is really signed or authenticated that way.

Another good idea was RFC 8509. And that's the, I think at the time we refer to it as a sentinel, but I'm not sure if that's the official term for it. But I think people involved here do know what that meant. You could basically ask a resolver what trust anchor it has configured, and it was either positive or negative. It either validates or doesn't validate with a trust anchor, if I'm not mistaken.

The issue there was you can only do that by querying resolvers and not all resolvers are opening to querying. So you need to either use something like the right NCC Atlas tool or the stuff that Geoff had been doing with an APNIC with Google, the Google approach. But at that time, 8509 wasn't really that much deployed. Now the risks that we identified in 2018 were twofold. The biggest risk was the revoking of the 2010 trust anchor, in the hope that everyone had configured the 2018 trust anchor. So only after the trust anchor was revoked from 2010, only then we could see if things were really properly configured.

And the other biggest risk we identified beforehand was response size increase. Let me start with that. Response size increase should not be an issue this time around, because we've been there. We've done that. We saw some issues, but nothing fundamental in order to stop the roll at the time. So now looking at the 22x, let's call it 2029 or 2028, the next rollover that we are doing. I'm sorry, I get two dates confused. So the key that has just been introduced in January, I think, the key that we generated last year in July. Like I said, the response size increase should not be an issue right now.

And then we have the other big issue is to see if people have configured the proper trust anchor, right, the new trust anchor. And we can only

see that if we revoke the 2018 trust anchor. Together with Paul Hoffman, we discussed this and we pushed back on monitoring with a purpose that if we monitor and results come back and then we have a threshold like we had last time around, like 0.5% or something around that, and that should then be a no-go. Anything based that's a no-go or go on monitoring that have just been described, that is the monitoring that is really, really inconclusive is not useful in our opinion.

Naturally, we will monitor. We've always been monitoring. There's ITHI data, there's magnitude data. And I can show you a little graph of them that I've produced. I'm sure that Duanne has similar graphs, but this is the 8145 data that we currently received, but this is only from L Root, not from the other root server operators. I'm going to share my screen. I'm not doing this often but let me just try that. I'm clicking on share and then Microsoft PowerPoint question. Go on. You can do it. Share.

GEOFF HUSTON: Success.

ROY ARENDS: Hey, thanks, thank you, Geoff. All right. So, this is what we see currently. This is sources over time by tag. Tag is the trust anchor signal that we see from resolvers. And this is not the number of queries we receive, but it's the number of hosts or the number of resolvers that are reporting it. So this is why you see a nice, I want to say diurnal, but it's a weak kernel pattern on the top line.

So you see five dots nicely aligned together and then two dots, which is a weekend, slightly below that. And what you see there, the top line, that's really the current trust anchor. That's TA4F66. I don't know what it is in decimal form, but I do know in hexadecimal form, that's the current trust anchor, the 2018/2017 key. You see another enormous amount of cruft there. Let me just do this.

And if I go to the next slide, now you see a lot of crufts removed, but still the cruft with 4F66 trust anchor in it. So you can see that even if a resolver configured the right trust anchor, it signals an enormous amount of stuff around it. We've seen this before. This should be nothing new. We've seen this before the last signaling. But if I filter out all of the unexpected ones, so we only look at 4F66 and the new trust anchor 9728 that has been created in July and introduced in January, we see something interesting. We see that it goes up slightly if you look at the orange line in the bottom right, it goes up slightly.

Now, I don't know if this is expected, but it's slightly surprising because a resolver, according to RFC 5011, has to see a trust anchor continuously for 30 days. We've introduced this key on January the 11th in the root zone of this year. And as you know, 30 days have not been passed but you still see an update. Now, I think this update is related to organizations that either ignore these 30 days with a broken 5011 implementation or they fetch the trust anchor from the root zone from the IANA site where we list the trust anchors. I see a hand up. I'm happy to hear you, Duane. Ask a question if you like.

DUANE WESSELS: Sure. Thanks, Roy. So yeah, I've been looking at this data as well. I'm pretty sure that it's from either software that fetches the root anchors on xml file or somebody manually configured it, manually looked at that file and configured the trust anchor in some set of software. But I had a question about something you said a while ago.

You were talking about revocation. And I think I heard you say that the point of revocation was the only point at which you could tell if a resolver had the right trust anchor or not. Did I hear that right?

ROY ARENDS: Yes. I might be using the wrong term here. It's not the point of revocation. I'm sorry. I mean the point where you stop using. Sorry, that's what I meant.

DUANE WESSELS: You meant just when the signing switches over to the new key.

ROY ARENDS: Yeah.

DUANE WESSELS: Okay. Thanks for that clarification Yeah, I'm done. Thanks.

ROY ARENDS: All right. And sorry for that. The good thing this is a bit more informal than planned. So you can see here from the monitoring that the new

trust anchor is being picked up. And as Duane said, this is likely due to people being proactive and software being proactive rather than 5011. Now, I expect to see this go up by 11th of February, which is the 30-day time. And it shouldn't go nicely up. It should go significantly up. But that's expected. And that will go up over time as well.

So that's the kind of monitoring we do, we've done in 2018. But I'm really, really, really hesitant to state something along the lines of based on the monitoring we see on either 8509 or 8145, and things like 0.5%, that's going to be, yeah, it's going to be really hard to defend given what happened in 2017. So I really urge the team that's responsible for the rollover not to be dependent on that kind of monitoring.

Now, of course, when something we know really bad is happening, like a known known, then we always have a plan for that, as you know. The rollover plan contains a part that's called the back out. But other than that, based on this kind of monitoring, this monitoring is nice to have, we welcome it. Also hope that other teams are going to monitor and share, and we will share of course as well. But to me, it's a nice to have, and really nothing more than that. All right, Geoff, go ahead.

GEOFF HUSTON: Yeah. Thanks, Roy. I appreciate we're running very low on time. Is there any plans to make this L root monitoring data public or is it internal to OCTO?

ROY ARENDS: Oh, currently is internal to OCTO, but I have no issue. This is this is really just another line in my magnitude scripts. I have no issue to publish this, but it will be OCTO, it will not be IANA. I don't want to bother that team with this stuff, but I'm happy to make this public. It will take a few days, if not weeks, because it's not currently on top of the list, but it shouldn't take too much time.

GEOFF HUSTON: Understood. Any other questions for Roy? Okay, I've given you enough time to figure out how to raise your hand or not, so I assume not. Thank you, Roy. That's very informative. I'm not sure on the APNIC side if we have any plans at the moment to resume measuring on 8509 but I will liaise with the team at IANA when we make a decision on that. And obviously when we do it, that data would be public. Okay, that brings--

ROY ARENDS: Well, can I just end with one little note? All of the stuff I've said here is basically my and Paul Hoffman's opinion on this. And it's completely independent of IANA. And this is just OCTO, if that makes any sense. And I've got a request from Kim to talk a little bit about the monitoring we have, which is what I've just done. So I hope this helps.

GEOFF HUSTON: Thank you. Thanks, Roy. In the time available, i.e. almost none, are there any other items of business? No.

DANIEL GLUCK: Do you mind if I pop in quickly?

GEOFF HUSTON: Absolutely.

DANIEL GLUCK: So following up on the email that I sent out on earlier this month to, I guess, go over last month's minutes and things like that, I had mentioned that the potential cadence of this meeting being the last Monday of every month that 1900 UTC, I haven't heard anything against that. So if it's okay with everyone here, it's okay with you, Geoff, I can get some calendar invites out for the next few months. And yeah, how does everyone feel about that? Any comment from Willem. I'm sorry?

GEOFF HUSTON: That's fine, I think. Like last year, if we don't have business, we won't have a meeting. But certainly, we've got something to discuss for the next meeting, the follow up to that report done by Willem and that crew. So there is something on for next meeting. So with that-- Sorry.

KALINA OSTALSKA: I was wondering if we could push by an hour. I think I put this in the notes somewhere maybe in the feedback form. In this hour, this particular hour of today, and in general, I have a weekly meeting of the team, so it's a bit of a conflict for me. I decided to skip the weekly meetings today, but going forward it would be really nice if we were able to push by an hour.

GEOFF HUSTON: Understood. Dan, can you do a doodle poll? And when you say push, an hour later or an hour earlier?

KALINA OSTALSKA: Later, sorry. Either way, but I feel that it would be better for American folks to have it later. Either way works for me.

GEOFF HUSTON: Thanks. Dan, if you could do a doodle poll as well. Thank you. Okay. I think that brings us to the end of the meeting. Thank you very much. Thank you, Willem. Thank you, Roy, for the material presented here. I think it's useful. And we have an agenda item for the next meeting to consider if there are any follow-up actions from that report that RZERC had asked for years ago to see if there's any further follow-up actions. With that, thank you very much, and we'll see you in a month. Thank you.

DANIEL GLUCK: Thanks, everyone. Bye-bye.

[END OF TRANSCRIPTION]