

JEFF OSBORN: Good afternoon. Good evening, good morning for those who are online. I'm Jeff. I'm going to invoke the new guy rule if I mess anything up. So, feel free to wave your hand. If we get anything wrong here. The agenda review. Require the agenda. And if there's nothing missing or nothing on there we don't need. Give you a minute. Then I'm going to pass this off to Dan for the caucus this engagement.

DANIELLE RUTHERFORD: Thanks, Jeff. And for anyone in the room, I've sent out around a few clipboards for attendance just for you to circle your name. There should be one in the back of the room. If you see me or see the one in the back of the room by the water table, fill it out, we'd love to give you credit for being here.

So just an announcement on the 2022 RSSAC Caucus member recognition awards. So, the RSSAC recently voted to recognize four RSSAC Caucus members for their participation and work in 2022. And I think our note went out this morning on the mailing list. You can see that Afifa Abbas, Anupam Agrawal, Ray Bellis and Amir Qayyum have all been recognized with high level of engagement and their outstanding contribution to the work of RSSAC Caucus.

Afifa served as the fellowship program mentor since 2021. During her service, she did a great job liaising between the mentoring committee and the RSSAC leadership in order to mentor the ICANN fellows in the best way possible. Anupam led the RSSAC047v2 work party that closed in 2022. As the work party leader, he contributed substantially to the publication of RSSAC047v2. Ray Bellis has been leading RSSAC002v5 work party that started in August of last year, and he was also a major contributor to the RSSAC047v2 work party. And Amir has served as the RSSAC liaison to the ICANN Nominating Committee in 2020, 2021 and 2022. In addition to offering significant volunteer time in NomCom, he's also served as an ICANN fellowship selection committee member from RSSAC since 2019. So, thank you, and congratulations!

Next up we'll have updates on work parties and work part work products. If you are giving any work party updates, we ask you to come up to the front of the room and this microphone. And

---

anyone who wants to get up and get up and ask questions, please join in in-room microphone.

First up, we have Ray Bellis.

RAY BELLIS:

Thank you, Danielle. This will be quite quick. So, RSSAC002v5, we've actually concluded the work party now, as of a few weeks ago. There really no major changes functionally RSSAC002v5. There have been no new metrics added or indeed no metrics removed. We had a quite extensive review of the existing metrics. Thanks to Paul Hoffman for herding the cats on that one. The, the main changes have been clarification that the traffic flows, we talked before about port 53 basically, they are only now related to port 53 traffic. Yeah. As a later date, if DoT, DoQ whatever become used on the route system, then we'll revisit that but at the moment we've clarified that... Yeah, we're only interested in the port DNS TCP 53 and DNS 53 traffic.

We've also made additionally the voluntary text, read me dot text that several operators are encouraged to put at the top of the RSSAC stats upload directory. But they've seen free form notes for the consumption of any user of the of the data. So, they can't stand any constraints on the data or any non-conformances, or indeed things like gaps in the data that might be there. Other than that, really not much has changed and thanks also to Andrew, and Ozan for the help for running the group for me. That's it.

Any questions? Anyone, questions? Okay.

PAUL HOFFMAN:

Hi there, I'm Paul Hoffman. So, this is sort of a second presentation. I gave one about a year ago. What we're talking about is RSSAC047. One of the recommendations was that there should be an initial implementation. I've sort of done that. And I'm going to describe this sort of stuff here. Next slide, please.

So, these are the words out of RSSAC047 about the initial implementation. Just to level set, this is not a full implementation. RSSAC did not say "please go and do the whole thing". It was like "please have an implementation available, so we sort of know what's going on". That's what the words are. And then the last sentence is the insights learned from this implementation will inform future revisions of this document. We've already had the

RSSAC047v2 document. There might be more later. But this is a way of starting to get data and stuff like that.

Next slide.

So, we've... As I said, at the last meeting, which was about a year ago, we've got something written. It's running. It mostly works but not completely works. And I'll have a couple of slides on that. Data is being collected not as completely or exactly in the same way that 047.. Oh, sorry. I was like I sound muffled, and there's a reason for that. Data is being collected now, even as we speak. And again, like in 047, it said, please do... you know, please run this every five minutes and such, that's happening. We did find some issues with 047 that we will probably want to discuss in the Caucus at a later time. When we did 047v2, it was just to correct errors. And so we limited it to that.

Now that we are starting to collect data, and people can say, oh, this is what it's looking like, we might have a discussion later. We might not. It might be that we say, okay, this is sort of good enough, we'll leave at this. But to be clear, the initial implementation that I'm talking about meets, does not meet all of the requirements. But it does show that the model that we set out in 047 is a viable model at least for collecting data. Whether the analysis is the right thing we want to do is completely different. This implementation is just saying, can we collect things sort of like we described and the answer is pretty much yes.

Next

So the in RSSAC047, it said, the vantage points will be distributed this way and such like that. That they need to be in major metropolitan areas. And the last bullet was sort of a future looking bullet again, this is all text from 047. That there might be a long term plan to actually acknowledge the fact that networks have the topology just geography, and so you know, RSSAC caucus might look into that. But this is sort of where 047 talks about specifically about vantage points and then next slide.

This is what we discovered in doing this. So remember there were those five regions. Turns out trying to get vantage points in South America, trying to get a bunch of vantage points in South

America is really hard to do. That there are only a couple of data centers other than everyone, everyone who has a South American presence basically is in Sao Paulo. There are some others they're not as stable, finding people who will let you have a VM in one of them is time dependent. Sometimes they're like, oh, no. We decided not to have those VMs anymore or well can I put an actual pizza box? And oh, no. That's like too complicated and such. So the initial implementation all know, has one in Sao Paulo also has one in Bogota, which is mostly stable, although it's V6 connectivity is probably a work in progress. And then one in Mexico City, which for those of you who know routing know that anything in Mexico City is basically in Texas.

Or, you know, it's closer to Texas than any like, Mexico City is closer to Dallas than Houston is in many ways. So, but at least we did that. And then in some of these cities, the connectivity is, even though these are all in data centers, it's somewhat spotty that is, we would not expect to see routing problems from a data center. As part of o four seven, remember we're doing probing but we also then do a trace route every five minutes as well. In case the probes go like, oh, someone so not doing something. We could then look at the trace route. The trace routes from some cities are from some of these data centers are surprising.

So that's good information. We can look at that. Unfortunately, all of that is separate. You know, all of the trace routes are being kept separately. And then there's the problem that sometimes vantage points fall over. You know, they're running just fine. You look at them, they're just fine. You can even log in to and say, why are you not, you know, passing out any data goes I don't know. I seem to be fine. So, this is something that you know, for a longer term delivered service, we're going to want to do some like, real monitoring and such like that because I've discovered a vantage point that is up, has not been actually transmitting any data for three months. In in some cases, it's because outbound rsync wasn't working all of a sudden. In some cases, because it has no data, it wasn't collecting it. But I mean, there was certainly one case where I was like, oh, we're not getting anything from here and it's been like a month. And I go what's wrong I just sort of kick it on side, and it's been collecting data, just hasn't been sending out.

So this is something where basically when a real implementation is done, probably going to want to have real monitoring on it. And but the flip side of all of these problems is it is actually very easy to add vantage points. And I noticed the other day, I forgot which, I don't remember who just opened up you know, allowed new vantage allowed new VMs in Honolulu. And like, okay. Great. We'll throw one in and it worked just fine like it came up such like that.

Next.

So the other problem that I've had is the correctness checking. It's easy to say everything everything's always going to be correct. But the way that we ended up wording this in RSSAC047 is that responses test against all the root zones that we're first seen in use in the forty-eight hours, preceding the query. And as we remember, remember we talked about this in the DC workshop, some queries are going to happen right around the time that the zone change, the root zone change. And some of the queries are very specific to the root zone. So it might actually be the root zone that was collected slightly in the future and such. I tried this in many different ways of, okay, I've got you know, a response here, and I've got this set of root zones for this link.

Let's make sure... And I still any up getting things that say, No. It's not correct. And we know it's correct. We know that the RSOs are doing correctly. So this... This was sort of a lack of creativity on my part. I'm pretty sure it's sol, but I spent some time bashing on it and I was just like, now, I'll wait a sec if I do this. So that's going to be something for the real implementation we either have to come up with a different rule or somebody who's better at figuring out how to juggle these two cues than I am, is going to have to be working. So right now, it's not doing correctness checking because we I mean, we absolutely believe everyone's being correct.

So next

I think this is the last slide. But so, in short, it's running now like I say, the vantage points are probing being still every five minutes. The central host has collected about fifty five gigabytes of stuff. I sent Dwayne like, a month's worth of a while back he looked at and such like that.

And again, this is just initial implementation. We can't like be using the dated to actually figure out who's doing what? But it is useful. I can certainly pull stuff together for anyone who's interested if they want like a bunch of the data and things like that. And right now I'm waiting for updates about you know, what vantage points basically, we really want. And if I can find someone more creative than I am doing.

And then I think maybe there was one more slide.

Oh, okay. Sorry. I should have included the Url for the Github repo. I've done that in the past happy to do it again. But this is all in the Github repo you all can look at it and like... I'm sorry. The code is in the github repo. The data is not, just because they would be unhappy to put gigabytes of stuff there. But I'm happy to send them around. So that's the status check for now. I'm not proposing we work on this now. But I'm also not proposing that we don't work on it. This is something for when people when we're thinking about what do we want to do to move this forward. As as a point of going. Okay. This is what we're doing. So Any quick questions now or else I would be happy to do stuff on the list, but Duane?

DUANE WESSELS: Thanks. This is Duane Wessels. Paul, can you, can you say a little bit more about... So, in RSSAC047 that has idea of threshold and monthly reports, I get the impression you're not doing any monthly reports because of you said the data's is a little bit you haven't ever...

PAUL HOFFMAN: Yes, I have... I have the structure of a monthly report, but because the data is so flaky, I'm just like not doing like, no. I don't want monthly reports with red blocks in it. Even because we know that the RSO are probably doing it correctly and these are routing and data anomalies not an indication. But, yeah, I wrote up, wrote up the code of how to take the data and turn it into a monthly.

DUANE WESSELS: I mean, I guess at some point, I think it would be useful if you know, even if the date is iffy and even if are some breads. I think it would still useful least for me to see, like, is what happens when we run our data through and generating report just to like... I mean, just as kind of another understanding.

---

PAUL HOFFMAN: Bad data, bad report, but that's still like, you know, I can put red at the top of the report instead...

DUANE WESSELS: Because for example if there was something like we got totally wrong on latency or availability. I it might show up in your reports even though it you know, anyway.

PAUL HOFFMAN: Sure. Okay. Yeah. This is thought. I can work on that a bit and make sure that those look sensible even if the answers aren't great.

ANAND RAJA: Hello, this Anand Raja. So, can we have some vantage points from kind of independent research? We can, and we can collect data.

PAUL HOFFMAN: Certainly, anyone can set up vantage points at any time. And so the code for it assumes that vantage points have been set up, it doesn't set them up. So, if you wanted to, for example, set up different vantage points and then run the code, you absolutely could.

ANAND RAJA: Okay. Yeah. And the central system, the vantage points will be communicating with the central system. Right?

PAUL HOFFMAN: Yeah. There... It's central system is pulling them via rsync every fifteen minutes.

ANAD RAJA: Okay.

PAUL HOFFMAN: Great. Okay. Thanks. And if again, if there's more interest in moving this forward, let's just talk about it on the list and such but, Duane I will follow up and spend some time on the reports that just look sort of sad. Even though they shouldn't. You know, like, the... They're clearly sad because of routing, not because of RSOs... Okay, Great. Thank you.

DANIELLE RUTHERFORD: I think next up is Dessalegn Yehuala.

DESSALEGN YEHUDALA: Hello. I'm Dessalegn Yehuala. The statement of work where proposing is about enhancing the metrics usability, existing Root Server System metrics usability setup. It's proposed under the assumption that existing metrics in their current form are only capable of measuring trends.

---

So, in order to have a broader picture of the root server systems, having additional models like a baseline model we're having some elements which can characterize this dysfunctional behavior of this infrastructure is important. So, I don't have much to say. That's essentially about extending existing metrics capability to one step up if Caucus wants to move this forward. It's open for discussion.

KEN RENARD: This is Ken Renard. Are there any particular metrics that you were looking at? You know, some of the metrics like RSSAC002 are just behaviors of users of the roots server system. Anything in particular that you're looking to tease out and then analyze.

DESSALEGN YEHUALA: The assumption is the current, the metrics in their current format only good in enough for measuring trends. We want to have a model which has some bare minimum, which to characterize the dysfunctional behavior. So probably, it could be seen as a revision of existing metrics. Probably that may result in additional metrics or enhancing the existing metrics. So I don't exactly know where this will [INDISCERNABLE], but it can be seen in both direction, revision of existing metrics.

If the Caucus feels appropriate introducing additional metrics, maybe also this is end result of this proposal, statement of work.

PAUL HOFFMAN: This is Paul Hoffman. So, we have two sets of metrics that are currently being collected. 002, which are self-reported and 047 which are coming from outside and doing that. For this work, you said we might add new metrics. Do you have a feeling if they would be new self-reported or new outside ones? Or are we waiting until the work has been done more to know what kind of metrics would then feed into that kind of anomaly detection.

DESSALEGN YEHUALA: My view, at this early stage you know, there will not be any alteration of existing operational model. We can use existing metrics by enhancing them to enable individual Root Server System operators to self-certify the heads of their infrastructure.

PAUL HOFFMAN: Okay. Very good. Thank you.

RAY BELLIS: So, is there a public draft of the statement of work available anywhere yet?



---

DESSALEGN YEHUALA: Yes. We have already produced the draft.

RAY BELLIS: Okay. I haven't seen any sign of this work. And I just looked in the caucus archive or mails that I've received. I haven't seen any sign of anything relating to this. Yeah.

Ken stressing, maybe Steve can someone out. But yeah. This is the first I've heard of this work, which surprises me.

DESSALEGN YEHUALA: Yeah. So, we have already produced a draft document. Probably we mentioned that anytime soon it's already it's the draft document produced. So the document produced between two ICANN staff members sending me, it may lack some detail, but we already have, I think this this statement of work, comes from my suggestion in RSSAC survey conducted a year ago and the ICANN personnel contacted me to flesh out what I said in the survey and then [INDISCERNABLE] we are ready to we can share it and then we can evolve it to the next level.

RAY BELLIS: Sorry. Ray Bellis again. Sorry. Having not seen statement of work, can you describe any proposed metrics that you're talking about because as we've discussed, we've already got..

DESSALEGN YEHUALA: So, we can talk about the latency existing metrics.

RAY BELLIS: Yeah. I'm just trying to understand the scope because we already have metrics to look at the trends in the DNS, which is RSSAC002 typically. And we've got trend, well we have metrics as we're looking at see correctness performance with RSSAC047. I'm really trying to understand what gap you're trying to fill. And I recall... Yeah, we had some discussion in RSSAC002 from suggestions of yours of metrics that were not in scope. So, I am trying to understand what you're trying sit, where are trying to sit that isn't already addressed by RSSAC002 or RSSAC047 or any other existing RSSAC documents.

DESSALEGN YEHUALA: I'm talking about baseline models. Current metrics as they stand, some of them have some threshold values. But we want to have the threshold values probably might only be helpful for no engaging the trends. So, this baseline model may come up with some statistical machine learning algorithms which can try to, at least, help the root server system operators to figure out outliers.

So, we're going have some either... I don't know. It's early stage to talk about probably maintenance be appropriate this is at this stage. We can talk some, for example, for some of the metrics, we may have to revise and include quadrants. I haven't seen the the existing, probably you may correct me if I am wrong. I have since in some of the metrics is proposal like ninety-five percentile and so on. But I don't see an exhaustive threshold values in terms of either statistical or machine learning. That can be expressed.

PAUL HOFFMAN: So, this is Paul Hoffman. I'm sort of answering Ray here but I am echoing a problem that we had when we wrote 047, which is terminology. There are measurements and there are metrics. I believe what you've said earlier is you're not proposing any new measurements but there might be new metrics that come out of this not based on trend analysis, but looking for specific anomalies. Is that correct?

DESSALEGN YEHUALA: Yes.

PAUL HOFFMAN: So this would really be new anomaly detection metrics without.. And so far, no addition of any new measurements. Is that? Yeah. So that sounds interesting to me in the sense, not that I am a data scientist at all, but it sounds interesting to me of looking for anomalies from current existing data in a way we didn't. That doesn't actually change anything especially in 002. Because 002 is, when it was originally, the original idea was hey, they're going to be adding the zillion new TLDs, they being ICANN. Sorry. I'm on staff at ICANN. So, I'm one of the they. Fortunately, I didn't add any of them.

And therefore, we should be collecting data as we can to look for stuff. And now you're thinking of a new use for data of, it would be nice to know if something broke even that's not related to the size of the root zone. So, I think this is interesting. I hope we don't get into a situation where ad hoc new measurements are proposed. Because it turns out, those are a little bit more difficult. I'm not saying it shouldn't. But I would love to see this start with a, given this data set what other metrics can we have? So, I look forward to seeing this.

RAY BELLIS: Ray Bellis here. To slightly return, turn the phrase appropriate use, I would like... I would not like to see arbitrary or ad hoc new metrics introduced without good reason. Yep. The RSOs are not

required to meet any specific metrics at this point. I was thinking of that with my RSO hat on. I will resist any attempt to come up with metrics by which you are attempting to measure an RSO, that has not been agreed at the very highest levels.

DESSALEGN YEHUALA: Okay.

WES HARDAKER: Wes Hardaker, ISI. I would suggest that to go forward, we need to see the statement of work. I mean, I have thoughts, you know, similar to Ray and Paul and you know, without reading sort of what you're proposing, it's hard to have the discussion. So, I would say at this point, submit it to the caucus mailing list, and then we'll discuss it as a caucus, but there is a... It's extensively expensive to produce RSSAC002 metrics and everything else and even 047, Paul has put it in a huge amount of effort into doing that. So, we do have to be careful about costs associated with that. Not to say this is not worth doing. I do look forward to reading it too. So, I think you may be proposing interesting things. But to discuss it, we need to read it.

PAUL HOFFMAN: Paul Hoffman again. So, Wes just brought up a really interesting thing that I had missed, which was I'm assuming you will be running analysis on existing data you're not expecting the data producer namely the RSOs to be running the analysis. Is that correct?

That you will be... Or that somebody will be looking at already reported data looking for this. Is that... Was that... Okay. Great. I think that will be helpful as well. Thanks.

DESSALEGN YEHUALA: So, no more questions? Thank you.

DANIELLE RUTHERFORD: Alright. Next, we'll have an update on the document repository from Hafiz who is a remote participant today.

HAFIZ FAROOQ: Hello everyone. Can you hear me?

DANIELLE RUTHERFORD: Yes. We can hear you.

HAFIZ FAROOQ: Okay. Great. Let me share my screen. I hope you can see my screen now.

DANIELLE RUTHERFORD: Yes. We see a screen for Zotero, research repository

management.

HAFIZ FAROOQ:

Okay. Great. So, thank you very much, guys. This is Hafiz Farooq. I'm new member to the RSSAC Caucus, recently joined last year. And today, I'm gonna give you an update about the document repository system. And we are specifically gonna talk about Zotero platform which is a very nice research repository management software. The software was recommended by Andrew, myself and Baojun. We worked on exploring and evaluating the product.

So a little bit about myself. I'm Hafiz Farooq currently working at Aramco. Since last seven years before that I have been working in some of the telcos, so I work in the area of network and cybersecurity.

So, let's move into Zotero. Zotero is actually an open source software for citation and reference management. And developed by Global Community. You can install it on Windows, a mac or LINUX operating system? Zotero also has connector for different web browsers which you can use for auto generating citations foot notes and bibliographies. You can also install plugins in Microsoft if we want to use this data for research paper writing, which I'm gonna explain later. Also very important feature for Zotero collaboration, you can create your research collections and if you are a group of people working on some research, you can share the research material between the team.

Also, it supports syncing in the cloud. So, once you enable syncing in the cloud, then you can share the data across and it also supports the backup and restore. Zotero is a free software. With the free software, you have 300MBs of research space. But if you want to have more documents to save then you need actually yearly subscription.

So now, I will give you a brief demo of the of the software. Before I move to the demo, now I want to show you the Zotero website. So first of all, once you want to use it or you need to download it from the website and install in your machine. It has a very nice interface as you can see on the screen. On the left-hand side, you can see different collections and you can organize them nicely in the form of a different sub collection also. Then you have a research paper you which you want to add into your repository.

On the right side, you are seeing the metadata data. Any kind of metadata for any research paper you want to add or modify you can add them quite easily. I'm gonna explain practically also.

So once you install it on your machine, the next thing you need to do is you need to create an account on the website. And once you create an account, you can link this account with your software. And another the important thing is once you have an account on the total website, you can create groups and groups are actually created for sharing the search across the team members. For example, this particular group RSSAC Caucus was created by Andrew. And he gave me permission to read and write into this particular repository. As you can see on the right side, this particular group is actually owned by Andrew. And it is public open for membership and this is the access I have.

So, you can actually add lots of member into the into the group and they can share, they can modify, and they can view the search repository. So coming back to the software, this is the the software interface which I installed on my machine. On the left hand side, as you can see, you can see my library on the top. This is my local library which is now shared in the cloud. And the RSSAC Caucus is the one which is shared across the team.

We added some research paper so far into this recommended by the Caucus team members. And we are still increasing it. We're still adding more data into this. Adding the data, adding research paper put into this library is quite simple. You can manually add some of the search paper presentation, general article or any other type of document you want to add. You can just click on this green button on the top and you can add anything manually. Once you try to add something, you can populate these fields, this metadata field on the right side. Manually. And the other thing you can do is, you can add some custom notes if you want to add some notes into your research paper or something you can add. Another important thing is you can add some tags based on your research keywords. And also, you can link the research paper together here in the related tab, if you have any existing paper which you want to link together, so this is how you can link your research documents with each other.

So this is, like, manually adding the documents into the repository. Another nice way is that you can automatically add

some of the research document into repository using your browser. Let me explain it to you. So in in my browser, I'm using Chrome at the moment. I have a Zotero connector. So, once you have this extension, this connector installed in your browser, you can add any search paper, quite interactive, any document online, you can add to the repository.

For example, I'm showing you one of the, RFC9364 written by Paul Hoffman. So, if you want to add it to the repository, you just need to click on this button. And it is linked with your local repository and your also group repository also. So you just need to select the right collection and then you just need to press the “done” and it will be added automatically into the repository. As you can see, RFC9364 is now visible here. And if I click on the info tab here, it has automatically populated the different shields. And I have some tags. For example, this is related to DNSSEC, so I can specify some more tags. Cybersecurity, for example. So this is how you can add some tags or anything you want to add to the document itself. So one more thing is you can also add some attachments. If you want to attach some documents also, you can also do that. You have some other material which you want to add to your particular research paper. It is all doable but obviously, it is gonna take some space out of your 300 MBs space.

Another important feature is the tags. On the left-hand side bottom, you see these different types of tags. So if you wanna search any sort of documents like caching, for example, you click on it, and it just gonna apply a filter on your repository and you can filter them quite fast. This is one of the way you search the data. So This is a brief overview how we add the data into the repository. Now if you want to use it, use this repository for making some research papers. That is also I'm gonna explain right now.

Right now, you can see my Microsoft word. And I have added actually the Zotero plugin into Microsoft world. And if you want to add one of the research paper into your research document, you just need to click “add citation”, it will show you one of the search box like this. And then you can say, for example, DNSSEC, like if you are looking for that particular RFC. Yeah. So you can just select like that and press the button, which is gonna add the

---

reference at the bottom with full details automatically. And you can refer it in the document also. So this is a very easy way of building a centralized repository for all your research documents. And which will be accessible to all the Caucus members. And you can use it in writing your future research document and research applications.

So that's it. It's just a brief overview. I'm now open for any feedbacks any comments. Thank you very much.

DANIELLE RUTHERFORD: Thank you Hafiz. It does not look like we have any questions or comments in the room.

HAFIZ FAROOQ: Okay. Okay. Great. Thank you very much.

DANIELLE RUTHERFORD: I believe next up we'll have a discussion led by Duane Wessels on RSSAC001v2.

DUANE WESSELS: So hi, I'm Duane Wessels. I was a work party leader for an update to the RSSAC001 document, which is.. Ken doesn't like my mask. So RSSAC001 specifies expectations of root server operators. The work party recently completed its work on updating that document. And the expectation is that that will get published pretty soon. However, the work party will continue to work on updates to RFC 7720. In the 001 document, we now refer to it as BCP40 so I should try to use that terminology because that's the sort of the stable reference. And so RFC 7720bis is an internet draft which has been posted by Wes and Mark. With some potential updates to the protocol requirements for roots of operators.

So, if you read the current RFC 7720 document, you'll see that it contains exclusively "must" level requirements. And there have been some proposals on the Caucus list for adding some new code requirements to BCP 40, which might not be "must" level requirements. One of those refers to the use of the NSID EDNS(0) option. And the other is about use of DNS cookies. So those are two specific items that have been proposed for adding to the RFC, somebody might have others that they would like to also suggest. But at this point we have sort of those two on the table. And if there any questions about 001 or 7720, at this point, I can take those or we could just open it up to the floor more

generally, to the two proposed items to see if anyone has thoughts on adding those to the to the RFC.

I will say that I believe the proposals for, well so I think we have specific text and proposed for one of those, but maybe not the other at this point. Is that right Paul? Just the NSID.

So, there's specific language being proposed for the NSID requirement but nothing but forth for the cookies. And the specific text that has been proposed, I believe it uses “should” as the requirement level. So this would be a slight change, like I said to the RFC because it's all must. I think a “should” requirement is a little bit new to that document and may require adding some additional text about, you know, why we have sort of both or mixed levels of requirements I guess.

They're negotiating who's gonna... He's gonna go first.

WES HARDAKER:

Yeah. So, this is Wes Hardaker and the current RFC 7720bis draft is... it was published by me, but going back to history of it. Right? So there was originally RFC two something mumble mumble.. It was always designed to be the minimum level expectations in the same way that within RSSAC, we have been talking about the importance of you know, RSO independence and we each bring different things to the table and different. So what is the minimum set of DNS requirements that we should all take on? And the latest, you know, series of potential changes bring about some recommendations rather than requirements. Right? So it's recommended slightly above the minimum. And Paul nicely, you know, submitted text in a in a... which I haven't merged yet because we haven't come to agreement. This is really what I wanna ask, right?

For me, question one starts with “do we want to change, you know, the purpose of this document to include recommendations understanding that, you know, they don't that RSOs don't have to implement them if they're recommendations”. Right?

Do we wanna change that and add things like you know, NSID and cookies or, you know, anything else that we might come up with. The other reason for making 7720bis was to change the reference to 001 because it said, you know, a specific version, and I made all the change... All the wording changes just to say, 001v2 or later, and... But we can't say BCP40 like we could. Yeah. You know, on the other direction.



So, I changed it, you know, So... But there's been some argument that we don't really need to do that unless, you know, if that was the only reason for publishing a new RFC, it's not worth sending to the RFC editor and everything like that. So that's where we stand. But I would love to get an opinion on you know, it's different time. Does the RSSAC Caucus believe that, you know, it would be beneficial to put in some of these other recommendations? That was not as distinct as I was hoping. Sorry.

DUANE WESSELS: No, that's that was great. Thanks.

PAUL HOFFMAN: Welcome to the world of Jet lag. So, this is Paul Hoffman. So, I'm one of the people pushing for adding “should” level for one thing that RFC 7720 doesn't sound like it's only talking about the bare minimum. Talking about what is needed to do operations. And we here in this room might know that really means bare minimums and such like that. But that's not the way it reads. And it's also not an RSSAC or, you know, an RSSAC document. This is a document in the IETF that was originally produced via somewhat the IAB and such like that. And the people reading it are completely different. So they may not understand that.

Then the other reason why I thought it would be good to have “should” level, which is not just you might do it. But anything that's a “should” level needs to if you're actually following the rules for us should, “should” level requirement needs to have wording about when it's okay not to. And the first one that we brought up was NSID, the world will not fall apart if some root server operator stops doing NSID, stops announcing it. Or announces it with non unique addresses.

DUANE WESSELS: I think your text has debugging in there.

PAUL HOFFMAN: Yeah, which is sort of the normal reason why it's there. So, I would be extremely hesitant to say must do NSID with unique things. Because if someone does it wrong, it's not going to affect almost anything. So, I don't want, like, add a bunch of “musts” but I also don't want to leave it out. Because it is useful as something that the root server operators might do. And therefore, that might come in. So that's why I started with NSID, cookies -for those who you didn't hear earlier as I was talking to Duane- Duane calls cookies the things that we shall not call cookie, you know. Right.

That may be something that is not useful. Cookies are in some people's minds useful for operations when you're under attack and such like that. That would take more explanation.

So, but again, the world is working fine today with some root server operators not doing cookies. Therefore, I don't think that putting a “must” on is going to be the right way to go. And therefore, we have this “should” thing where it's “should” with an explanation. So that was my reasoning, Wes, of starting there. There might be new “musts”. But I don't feel either of these are must level.

RAY BELLIS:

Ray Bellis, [INDISCERNABLE] hat on. I feel [INDISCERNABLE] cookies that even “should” is too strong. I would say, “recommended”, I think would be a perfectly good word for that. Trying to remember that. [INDISCERNABLE] 2119. Yeah. Okay. I'd like something a little bit lower than that sense that I think yes, they're a good idea, but actually, you know, Cloudflare nodes the operator for on our behalf by cloudflare, they don't do cookies at all. The world is fine with that. That works. Yeah. The ones that I run, they do. Cloudflare ones don't. Yeah. Big deal. But if we change that to high conformance level, then I've got to argue. Well, why don't Cloudflare do it? That's actually get a reason for it. It's just... Yeah. More and my feel was necessary.

WES HARDAKER:

So for, this is Wes again. If we're gonna dive down the individual rabbit holes, the cookie one, I actually don't understand quite as much because that's something an operator. In order for me to defend against the Ddos attack, I wanna implement cookies. Why would you care if I do? Right? Unless I fall over. I mean, it might be good recommended practice through operators you know, that you should do that. But why is it a requirement for me? It's so... I mean, cloudflare has other mechanisms for defending themselves against Ddos attacks they don't need cookies. So it's sort of weird to even make some statement of.. Right. And it's just weird to make a statement of “this is the technology you should use to combat this problem” when that is not necessarily the only one.

PAUL HOFFMAN:

So this is Paul Hoffman. We're now arguing about something for which there is no wording. So that's dangerous. But you asked the question, why should we even bring it up, you know, to that level? Many things that are currently in BCP40 are also not

technically needed for the root service, but the IETF together with the root server operators felt that it was. So, for example, IPv6. If a root server operator had a religious reason for not doing IPv6, as long as it wasn't all the root server operators, no one would notice. Or they said I'm doing IPv6 and somehow it stopped, no one would notice. It's the same argument. And yet, that's a must level there, for a reason that many people [INDISCERNABLE]. Agree. So if we look at cookies, I think we need to only look at a not say just say cookies, but especially if it's gonna be a "should", flesh out the arguments. And I don't wanna do that here because I actually don't think cookies necessarily is going to make the grade. But I want anything that goes on this list. Because it should level to be more than one sentence. It has to be two or more sentences. Saying, why is this here? Why is it okay not too?

WES HARDAKER: So if I can take your hypothetical bit further, if all thirteen RSOs stop serving IPv6, you would notice. If all thirteen stop supporting cookies, nobody would notice.

PAUL HOFFMAN: Right? But that's not the way the BCP40 is currently worded. If we want to reword BCP40, from the ground about the whole service, we can, but I don't think that that's a good use of BCP40. So again, the hypotheticals are what do we believe the root server operator should do within the structure of the current BCP40.

DUANE WESSELS: So I think as Wes said, the first question to answer is whether or not BCP40 should set the minimum or a little bit higher and since we've spent all of our time talking about specific things and nobody argued to keep it at the minimum, I'm kind of guessing that sort of the group feels that that it's appropriate to, for the RFC to go a little bit above the minimum.

Does anyone disagree with that? I mean, is there anyone that wants argue that the RFC should stay just at the minimum? Mark Blanchet? No, he's not... Yeah. I don't I know, I don't think he's online, is he? Yeah. But so maybe that's a discussion we should take back to the to the Caucus list and settle first and then if we have consensus on that, then we can move on to the specific items of cookies and NSID if that's good for people.

Robert, you get the last word. I think. Let's see. Let's see what you say actually.

---

ROBERT CAROLINA: Yeah, Rob Carolina ISC and RSSAC. Forgive me because I... I'm not intimately familiar with RSSAC001, which is to say I don't understand that at all. But my question is and this seems to be recurring question a lot of discussions over the last few months. Is this attempting to explain what every root server operator is meant be doing individually or is this an effort to communicate to the world what the RSS collectively will achieve?

DUANE WESSELS: RSSAC001 is addressing root server operators individually. And I believe the RFC takes the same approach. It does not describe the collective RSS.

ROBERT CAROLINA: That's my work done.

PAUL HOFFMAN: I'm glad you brought that up because I really think we have wording Problems. The abstract says the DNS root service, root name service is a critical part of the Internet architecture the protocol and deployment requirements for the DNS root name service are defined in this document. But then there are things about root servers and service and stuff. It's not clear. And I if we're going to delve in, I think clarifying it would be good.

DUANE WESSELS: Right? But wouldn't you agree recall that the intention is that the requirements in the RFC are two individual RSOs. Or you're ..

PAUL HOFFMAN: I don't get to say that. We have words in something that is labeled to BCP.

DUANE WESSELS: Okay.

ROBERT CAROLINA: And I guess, Rob Carolina again, I guess my last thought on this would be if the goal becomes is or becomes to document what the root server system, the RSS service collectively achieves, the first question I have is: Is there any rationale that falls off the back of that every root server operator is expected to do the same thing all the time? So if there it will, well that's an open question.

Well, it's because.. the reason that I ask is because I'm hearing a lot of discussion about how root server operators do different things and heterogeneity is part of the resilience of the system, etcetera, etcetera.. It seems to me that there's a there must be a

---

small crunchy small core things that every root server operator must do and must not deviate from. But there might also be a large number of other things which they could do wildly differently and the service as a whole would not.

DUANE WESSELS: So, we are running short on time and the queue is getting longer. Do we wanna cut off now so Brad has time? Yeah. Let's.. I'm sorry. Rob opened the can of worms. He knew he was doing it right. But we're out of time for this.

WES HARDAKER: When clarifying.. I am the current author of the current.. I did not author the last two. I just took the pen to help edit.

JEFF OSBORN: All right. Next, ironically, would be the GWG update. and I don't see Brad online. I guess, if there's, what we do in that case?

DANIELLE RUTHERFORD: It's just that, it's just an open information if there's any..

JEFF OSBORN: Anything with that. I'm... I'm about to start juggling so we can avoid 7720.

DANIELLE RUTHERFORD: If there's no meaningful substantive updates, I think we can move on.

JEFF OSBORN: Okay. Then I believe it's your turn on strategic trends.

DANIELLE RUTHERFORD: This is just an announcement that ICANN has begun the fiscal year twenty-five planning process and this includes strategic outlook identification sessions with all of the SO and AC community groups with within ICANN.

So, the RSSAC caucus will be seeing probably a doodle pull for availability between April and May. We really encourage a lot of participation among the SOs and ACs. This is a chance where the people setting the priorities for the strategic plan and the operating initiatives for ICANN org community get to hear directly from the community group. So if you want your voice to be heard as far as what are like the top priorities and what ICANN use be focusing on, we really encourage RSSAC participation in that process.

JEFF OSBORN: Thanks Danielle. And with.. oops?

PAUL HOFFMAN: I just clarify one thing is that people would be in this as individuals, not as representing RSSAC. Is that correct?

DANIELLE RUTHERFORD: That's correct. The identification session would be taking place, the planning team with host one for RSSAC, but you participate as an individual with your own ideas and your own preferences.

PAUL HOFFMAN: And the reason I jumped up about this is one of the things that comes out of this is.. Oh, it seems like people with disparate identification care about this new topic. It should be on the radar for ICANN. That doesn't have to be RSSAC cares. In fact, it shouldn't be. This is a way to collect individual voices.

DANIELLE RUTHERFORD: I would agree with that. Wes, do you have a comment?

WES HARDAKER: I do. Wes Hardaker. ICANN Board hat on this time. This is incredibly important. Right? ICANN does what the community wants it to do and setting the objective for what should happen in the future is up to the community at large. It's very important that we get as much feedback as possible.

JEFF OSBORN: Thanks Duane. And with one of the two remaining minutes, Ray?

RAY BELLIS: Static. Okay. Just a quick announcement for those of you who haven't already seen this. [INDISCERNABLE] around last week at Cancun. I've just released that to the ios and macos app store called DNS Root Explorer. Which they, little app that she shows a fairly decent map of the root system, do a country and city search of where root server instances are. Run a test at from your clients, and it'll tell you which anycast instances you are actually hitting from whichever network you are on at that time. Not actually trying out here, but anybody who wants to demo, I've got it here my ipad laptop. Happy to show it to anybody. Thank you.

JEFF OSBORN: Thanks Ray. The last item is our next meeting. It will be in person at ICANN78 in Hamburg, Germany. It will not be from the twenty first to twenty sixth. Do you know when the actual Caucus meeting is?

DANIELLE RUTHERFORD: I don't think we have that scheduled yet. That's just the dates of ICANN78.

JEFF OSBORN:

Okay. One of those days in there then. And with no further due, and exactly on time. We give the evening back to you. Thank you all for coming.