

---

BRAD VERD: ...from there. Tripti, is there anything more you want to add to this? No? Okay.

All right. I'm going to turn it over to Tripti, who's going to talk about the workshop.

TRIPTI SINHA: So close to a year ago now, the RRSAC, as you know, has most of its meetings at ICANN meetings. And there were many complex issues that we needed to deliberate. And we thought it was probably good to start having what we call workshops, where we essentially sequester ourselves for a couple of days in a room and answer some rather difficult and emerging questions.

So we had our first workshop. I believe it was last September. And we held our second workshop this past May. And so I'll give you a quick synopsis on the workshop. It's essentially a report. It has been released. It's been made public. And we held a workshop May 11<sup>th</sup> through 12<sup>th</sup>. It was hosted by Verisign. And we essentially took the approach of taking our last topic and delving deeper. So the metaphor that I use is peeling the onion. So with every workshop, we'll continue to peel the onion until we get to the bottom of answering the questions that we're trying to answer.

And it essentially centered around the same questions that were addressed in the first workshop, which centers around accountability, continuity, and evolution. It's in a manner where root server operations

---

*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.*

---

rest today, we need to now look at what's coming next and how do we evolve, so on and so forth.

So the way these workshops are conducted is we assemble a planning committee from the RRSAC. So there's a small group of six of us, your co-Chairs and four other members of RRSAC, that begin work roughly a month or so before the workshop; put together the agenda and the content. And then we go to our destination, and we answer the questions.

So we took the approach of looking at three broad areas. And the main one was architecture, looking at the architecture of the root server system. Moved on to the topic of evolution, and then something that we call reinventing RRSAC.

So I'll start with the broad topic of architecture. And we peeled the onion by saying the DNS root server system, very reliable and very robust these past many, many years. And many things have contributed towards it. It's the highly distributed nature of the service, any casting, and just the diversity that comes with how the 12 organizations come together, all very different, very different manners of operating. And in fact, the more we delved into the discussion, we said, "You know, it's probably one of the original cloud services." Cloud services today define a service that's readily available and not necessarily hosted on your site. It's up there in the cloud somewhere. And we're probably one of the original first services, even before this moniker became commonplace.

So we decided that it was probably a good idea to document what makes the server so reliable and robust. And Wes will come up later and

---

give us a quick synopsis on our statement. So that is actually being documented and will be published shortly.

Moving on, we talked about technical risks and benefits to the root server system, should we lost a unique identifier, should we lose one of the letters. And we spent many hours discussing that. And then we took the approach of perhaps we're asking the wrong question. Maybe we should come at this from a different angle. And the question that we really should be asking ourselves is, what is the maximum latency a party should experience when it's transacting with the DNS root service?

So what we've done now is this overlaps with another piece of work that just went to the Caucus recently on what's the appropriate number of any [cast] instances. There's a tremendous amount of overlap with that statement. So we are hoping that the outcome of this will be provided soon. So rather than take the approach of, how many more letters do we need or how many fewer letters do we need, this is the approach we're going to take. And so at this time, that's the part that we are moving forward with.

And moving on to the third item within architecture was that we felt it was imperative that the root sever operators make a statement that talked about our response as an operator to provide complete and unmodified DNS responses using DNSSEC so that a client can cryptographically verify the response. So Liman will come up a little bit later and give you a statement that we have just released. Has it been made public yet? Do you know? It has been made public. And he will talk about our statement that we released on behalf of the operators.

---

Then we went on to discuss evolution. So the service has been reliable, resilient. What is the next big thing? And how do you set expectations for operators? What's the designation process, which is not defined yet? So we decided, let's start with the technical answer. What are the metrics against which we measure a root server operator? And considerable work was done in this area with RRSAC 001, so we've decided to use that as a starting point. Duane is leading the study. We will be doing some tremendous work on that with his group. And this will then go to the Caucus and others for further comment. So we'll start by defining a technical bar against which we can measure our performance.

And the last item we talked about was reinventing RRSAC. And it came to our attention that there was quite a bit of confusion. So RRSAC, the way it's done today, is an Advisory Committee to the ICANN Board and community. But there's no mode, per say, to reach out to the 12 RSOs. So we decided that since this is an organization that contains all operators and others who are associated with root zone management, that we would be the front door. In other words, we don't speak for all the organizations, but we will certainly take in a question and then pass it on to the operators. And actually, we currently have three different items that have come our way which we are using this process to communicate down to the operators.

So that essentially is a synopsis of what we did during those two days. And we concluded that two days wasn't enough. So we're having our next workshop in October of this year, and it is now extended by one day. And we'll continue peeling the onion, go on to the next layer. So with that said, do you have any questions? If not, I'm going to turn it

---

over to others, who will give a little bit more detail on the different items.

So any questions, first?

Okay, hearing none, the next person on the agenda, I believe, is Wes. Wes, go ahead.

Go ahead. Let them know it's not fully baked. Yeah.

WES HARDAKER:

Yeah. One of the things that we talked about, as was said, was a statement regarding the impact of the unavailability of a single root sever within the root sever letters. And in particular, the statement, which I'm going to read, is, "Document to core underlying reasons why an outage or otherwise unavailability of a service of a single root server is not, and will not, pose an immediate problem for the collective root server system or for the global Internet."

We had hoped to finish the output document from this by this meeting. I'll take the blame for that. I didn't push people hard enough, and so there's been a little bit of negotiation about the wording that didn't quite make it to be finished by this meeting. So that will be coming out soon and will be sent to the Caucus for review and comment.

It's worth noting that we referenced a couple of important events and studies, including both the attack that happened in June, as well as the statement by the RSOs for both – excuse me – June and last November. And then, of course, RRSAC 003's list of TTL discussions and the fact that, actually, a lot of clients ignore the TTLs as sort of venues for

---

information. But a lot more information could be gathered, and probably future work by the Caucus could go into studying the problem at further length than we really had data to conclude with. But that should be coming out, probably in the next couple of weeks, to the Caucus for review.

BRAD VERD:

Any questions or comments to Wes or anybody else working on it?

No?

LARS-JOHAN LIMAN:

I guess it's my turn next time. We've also issued a document with a... I need new eyes. The RRSAC statement on client-side reliability of root DNS data. And this has the basis that we've seen rumors, or we've received input that indicate rumors, that people think that root server operators provide different answers to people for different clients. That is not true. So this is actually a statement from the root server operators, channeled through RRSAC, as that is a publication mechanism that is useful.

And the statement is just only half a page. Five points, so to speak. One where we assert we use the same source data, and that is the data that IANA provides to us, through the channel that you all know and love. We also make a statement that we support the [IEB] statement on the single DNS root on the public Internet. And we go to the extent of quoting that RFC. And we assert that we serve all clients on equal basis and we don't provide any difference in the answers that we give out.

---

And in addition, there is DNSSEC as a means to validate content so that people can actually validate that we don't modify the data, or that they receive unmodified data. And the final statement is that we, as root server operators, disapprove of tampering with the data because, as you all know, even though we give out correct information from our servers, there is no – my mouth – there is no way that we can prevent people in the network part from modifying the data. And it's totally out of the control of any server operator to do that.

So that's basically the statement. It's there. It's on the webpage, the public page for statements and publications from RRSAC. If you have any questions or comments, I'm happy to take them. But this is more of a report on work that's happened recently.

Seeing none here, I'll give it back to you.

TRIPTI SINHA:

Now I'm going to turn it over to Duane to give an update on the work he's doing with his group to define the technical metrics against which operators would measure themselves.

DUANE WESSELS:

Thank you. Okay. So I'm here to talk about this document, which is sort of listing a number of technical elements against which potential new root server operators could be evaluated. So the idea behind this is that if at some point in time there is an opportunity for a new root server organization, that such an organization would be given some kind of request for qualifications. And the response to that would be evaluated

---

by the list that's in this document. So this document is not that request for qualifications. It's just ideas on how that response would be evaluated.

So it's divided into five sections, these expectations. The first is design, which means that the candidate root operator would describe their design for how they would operate their root.

And scroll down this part.

So in the design section, it's actually relatively short. We talked about just overall service design, very general. That whoever was doing this evaluation would look for certain general characteristics, such as locations of services, type of networking, peering versus transit, and so on. The other design point is service availability, which again at this part is very general, and we have more specific thoughts on that later in the document.

The next section talks about networking and experience. So in particular, we would expect that any candidate root operator would already have experience being a DNS operator, and they would be evaluated to that extent. We talk about security audits. So a candidate may be asked to provide security audit detail, showing that they adhere to best current practices and so on. Talk about addressing resources, that the candidate operator should have its own addressing resources, because it's assumed, for example, that Anycast would likely be used.

And I want to say that in this document, we're very careful to say that not all of these things are requirements. They're not things that must be present, but these are just things that would be evaluated. Some are

---

stronger than others. Some are “should have.” Some are “may have,” and so on.

Also under the networking section, we talk about peering. So if an operator would use peering, then they should have up-to-date peering contact information and entries in peering databases. They should have up-to-date information in address registries. There’s an item that talks about evaluating their internal zone distribution architecture.

The next section is all about diversity. And we go through these whole layers of diversity, from geographic diversity to network provider diversity, hardware diversity, software diversity, and personnel diversity – so not relying on a single person, for example.

The next section is all about documentation. So the expectation is that a candidate operator would provide certain documentation on maintenance procedures, disaster recovery, business continuity, and so on.

And then there’s a final section, which is miscellaneous, where we talk about data and measurements. So for example, a candidate operator may be asked to provide sample RRSAC 002 data to show that they’ve got that implemented and can be checked off. We describe the possibility of an evaluation period so that a candidate would be asked to operate the service for some amount of time before becoming official, perhaps. Talk about participation in all the groups that we’re familiar with, such as RRSAC itself, DNS-OARC, IETF. Talk about participating in [Diddle] collections and so on.

---

So this document is not complete. The idea is probably for us to do one more round of polishing and then come to the Caucus for more input and participation from the Caucus. Yeah. Questions about this?

Paul?

PAUL HOFFMAN: The title of the document in the agenda is different than what you've said.

DUANE WESSELS: Because this is RSS?

PAUL HOFFMAN: No, because what you said was this would be used for evaluating, say, a new operator.

DUANE WESSELS: Yeah.

PAUL HOFFMAN: And from the title I saw, I thought that these evaluations might be applied to current operators, as well. Has that been decided? Is that something that maybe will come to Caucus?

---

DUANE WESSELS: In the discussions we've had so far, we have not suggested that current operators would be evaluated against these requirements or these metrics.

PAUL HOFFMAN: Okay. Right. Although I do hear a comment behind me, which I think I would agree with, which is but that might become inevitable anyways.

DUANE WESSELS: It might be. It might be.

PAUL HOFFMAN: But for now, this is for looking at new folks only?

DUANE WESSELS: That's right.

TRIPTI SINHA: We're coming up with a set of metrics. Absolutely, they will be used to evaluate new ones. But it doesn't mean it will not be used to evaluate current.

PAUL HOFFMAN: But the design of them is for new ones?

---

TRIPTI SINHA: It's for an ideal operator. It's not so much for new. It's, what are the metrics?

PAUL HOFFMAN: Okay.

TRIPTI SINHA: What are the metrics? So we haven't decided how it's going to be used right now. We're focusing on just the technical metrics. Yeah.

PAUL HOFFMAN: Yeah, okay. Thanks.

DUANE WESSELS: Yeah. And just to be clear, this document is only about the technical metrics. There are other non-technical metrics that you may want an evaluation. But that's not in this document.

TRIPTI SINHA: Thank you.

PAUL HOFFMAN: Yes. Specifically, there's no framework for actually using the metric, right? So it's a list of metrics, with no ability to apply it to anything at the moment.

---

JAAP AKKERHUIS: Well, I have a more generic question about these statements. And that is, given that RRSAC is an ICANN community, and RRSAC Caucus is an ICANN community, are these going for the generic public comments, like all the other documents for all the other committees? The ICANN public comments, normally people give you room to [inaudible] public comments on the document. All the other Advisory Committees are doing it.

UNIDENTIFIED MALE: No. I think that that's an incorrect statement. SSAC, Warren can speak for SSAC. But also, RRSAC has made other statements that have not been out for public comment yet. So either everyone's doing it wrong, or maybe you're mistaken.

JAAP AKKERHUIS: Well, no, it's...

UNIDENTIFIED MALE: The intent of these documents is not for public statement. This is output from RRSAC, and ultimately the Caucus, which is advice from RRSAC to the Board and the community.

JAAP AKKERHUIS: The reason why I ask is because if you are going to get the Caucus meetings at ICANN, certainly we will be getting this question from other people. We should be prepared, within the frame of ICANN, to have this.

UNIDENTIFIED MALE: Agreed, we will be prepared. And then there is also an opportunity for, if other members want to contribute to the solution, they can join the Caucus.

UNIDENTIFIED MALE: I think I now understand what Jaap was asking, which is, if we publish it, we're going to get comments. Not the official ICANN public comment process, which is the whole it has to be published for 30 days before, blah, blah, blah, with that.

JAAP AKKERHUIS: Yeah, and you might have people who will be pushing for that.

ASHLEY HEINEMAN: Just to clarify again, to make sure I understand what Jaap is asking for, I think he might be referring to policy development processes. And before ICANN can adopt any policy development, there's a public comment process on that. But in terms of an Advisory Committee, I've never heard for public comments when advice coming out of an Advisory Committee.

JAAP AKKERHUIS: RRSAC put out yet another document. It always goes for public comment.

---

UNIDENTIFIED MALE: I think we have a terminology issue here. But, yes, RRSAC has published a statement. We just published the statement that Liman covered, right? And the public, when we review it at the next RRSAC meeting in the public sessions, there might be public comments. But there is no formal public comment period, where people submit comments, because there's no place to do that. I'm sure SSAC is the same way. I'm looking...

UNIDENTIFIED FEMALE: [inaudible] the same.

UNIDENTIFIED MALE: Okay.

TRIPTI SINHA: We should take this as an action item to get this clarified.

UNIDENTIFIED MALE: Andrew, can you [crosstalk]?

BRAD VERD: Jaap, if you have specific examples, can you send them to the Chairs? Because I don't see them. Or maybe send them to the list so that we know how to do a comparison. No, I'm not saying at the mic. You can say, "This URL is a public comment on one of them," and that would help us understand.

---

JAAP AKKERHUIS:                    If you go to icann.org, you see the list of closed public comments. You will see SSAC documents popping up there.

BRAD VERD:                        Russ, you had something to add?

RUSS MUNDY:                    Yeah. Let me try to clarify a little bit from the ICANN perspective, as somebody that's spent a lot of time doing things in various parts there too. And that is each of the entities that make up and participate in ICANN in a formal-ish sort of way, whether it's SSAC or ALAC or GAC or GNSO, they each have their own processes. And that's what the RRSAC 000 document writes down, is what the RRSAC processes are. Now, the SSAC processes are a little different, but they are also bound in a procedures manual.

And in terms of what some people mean when they say, "the ICANN public comment period," I think Ashley hit it pretty well, in terms of the whole policy development process part of the GNSO includes a timeframe and a structure for doing formal public comments.

Neither RRSAC nor SSAC have a comparable similar thing. There are public comments that are brought up and talked about and discussed after a document is published in accordance with procedures. So I think that the whole idea of if there is, or what does it amount to, when you say, "public comment period," needs to be done in accordance with the

---

RRSAC 000 or, if we're talking a different part of ICANN, what that part of ICANN does.

So I don't know if that helps clarify or not. That was what I was hoping to do. Thanks.

BRAD VERD: All right, great. We'll take this as an action item and try to come back for clarity. Thank you, Duane.

TRIPTI SINHA: Moving on to the next item... Any other final questions on this topic, on the workshop report?

Okay, Ashely?

ASHLEY HEINEMAN: This is about the workshop report, but also all the other statements that get finalized and published by RRSAC. I know they're posed on the ICANN webpage, but is there any other way that gets it out of the black hole? As someone who might be interested in these subjects and statements, do they have to just physically find it at the RRSAC page? Or do they somehow get set out other ways, either through – anything, actually. Not just Caucus though, but more broadly than that.

---

TRIPTI SINHA: No. I know the RRSAC Caucus gets a copy that ICANN staff sends to them, but I don't believe... There's no other modality to get it out to a broader community. It's just posted on the ICANN website.

UNIDENTIFIED MALE: Are you asking that it happen, or are you asking whether it happens?

ASHLEY HEINEMAN: I don't know if I have an opinion on all statements. But statements such as, "The root operator is not modifying the data," that kind of message that I think is good for public consumption for an audience that's broader than the Caucus. So I think, at least in certain cases – I don't know want to bind my hand to all cases – but if there was some way to get them broader visibility, I just thought that might be something good to look into.

UNIDENTIFIED MALE: That's where hopefully the Caucus takes on a little bit of responsibility there and continues to spread the word.

UNIDENTIFIED MALE: Or there are 12 root server operators. I think at least one person at every one – I shouldn't say everyone – at least one person at, at least, ten of them have a blog and things like that. I think it would be very appropriate for an individual root server operator, in a blog, to say, "Look, we were part of this statement that was everybody." That kind of thing gets out a lot better than anything that ICANN.

BRAD VERD: We're happy to take that back to the root operators, share that feedback.

Okay, that's it for the workshop. Let's go on to current work, work parties and work products that are currently underway. The first one is the Root Server System Naming Scheme, which – is John here? Yeah, there he is – John is going to give a quick update on, as he is the work party leader.

JOHN BOND: Yeah, so we're discussing whether to change the delegation of the root zone from root-servers.net to something else. A number of ideas have been discussed. The work party has been going on for some time now. It's the few issues mainly that it's sort of stalled, because Joe, the previous document lead, had to pull away.

We've had a reboot. It's been progressing pretty well. We have a lab setup now that has all six schemes so people can test them against [naught], NSD, and BIND, to see how the schemes look in reality. As for the document itself, we mainly need to do editorial updates to the appendix, which was posing the lab setup. So we just need to fill that in with some data points. And also some final editorial stuff on the technical conclusions. And finally, that the actual recommendation that we are going to make, I think there's still some fine details that we need to agree upon in the work party to decide which scheme we're going to recommend. And the reason that there's a little bit of confusion there is that there's some differences in the way that the authoritative servers

---

handle some of the naming schemes. We just need to analyze that and agree on a recommendation.

Any questions?

UNIDENTIFIED MALE: [inaudible]

JOHN BOND: Yeah. So the reason this came out to begin with was there was a potential for a vulnerability in the way the current delegation exists, which could have been solved. An easy way to solve that would have been signing root-servers.net. And because that question was raised, we decided to look at the bigger picture of, should we still use root-servers.net? Is there a better name that we could use? Do we even need to have a separate zone which is delegated and has these labels?

DUANE WESSELS: Hi, John. This may be putting you on the spot, but could you describe briefly what the candidate choices are and maybe which ones are in consideration, or not at all in consideration?

JOHN BOND: Yeah. So we've discussed six options. The first option is no change. The second option is to sign root-servers.net. The third option is to move the delegation into the root zone, so they won't be in a separate zone like they are now, with root-servers.net. And the fourth option would

---

be very similar to what we have now, except instead of root-servers.net, the name servers would go under a new TLD. As an example, .root-servers, and that would be signed. The fifth option, we're sort of going back to how the names were originally allocated, where every operator would manage their own name server and their own TLD. But when we considered that, we also took into consideration the name compression algorithm. So we would make use of that. And the final option was to have one delegation, an NSSet with one record, which had multiple A and [cord A] responses. So the delegation for the root server might be A-to-M dot-root servers. And when you look up the A or [cored] A, you would get 13 or 11 records back.

I think the two options that we're considering at the moment – and I see some people in the work party, so please feel free to correct me if I'm wrong – is the option where we essentially have something very similar to what we're doing now, but we change the zone that the servers exist into a TLD, like .root-servers, or to move the record into the root zone itself. And the tradeoffs there are to do with the response sizes you get back in the priming query and how the additional section is filled.

BRAD VERD:

All right. Thank you, John. Tripti, you want to talk about the history documents?

TRIPTI SINHA:

So the history document, I think the Caucus has seen this document. Just to give you a brief history of the genesis of this, it actually is a work

---

item that comes out of the RRSAC. And we did this as a prework for our first workshop. And one of the themes of the workshop was evolution. Where are we going to? How are we going to evolve? So we said to do that, you need to understand your history. So we decided to collect all this material, put it together. I think we've sent it over to the Caucus. And it's an amazing piece of work. It's come back, and I offered to take a clean sweep and write it as it's read with one voice, because many people contributed towards it.

So I just wanted to apologize. It's still with me. I haven't finished it. And I intended to complete editing the document. We're not changing the content, just the language and how it reads so that it reads as a narrative. And I hope to have it done sometime in the early fall, and that's the status. And my apologies again for not having done this sooner.

BRAD VERD:

If I could just add that the history document is also where some of the Naming Scheme Work Party came from. We started to document everything, and it made us question, root-servers.net came from a long time ago. Maybe it's time to just question everything. That was the original start from it. Is this still the right way to do it? And if not, how should we?

All right. Tripti?

TRIPTI SINHA:

No, back to you, if there are no questions.

---

UNIDENTIFIED MALE: [inaudible]

TRIPTI SINHA: No.

BRAD VERD: Oh, I'm sorry. So, yeah, upcoming work, which is the Anycast instance statement of work that has been sent to the Caucus, correct, Andrew? No, it has not? Okay.

So there is a statement of work that is being finalized and should be being sent to the Caucus in the very near future here. This is the one, it started out as a number of different efforts. And it turns out that we've kind of collapsed a couple of different questions into one. And one of the bullet points that we're trying to address came out of the workshop, which is the latency question. So this statement of work should be coming to the Caucus soon. We'll be looking for a work party leader, and creating a work party to work on this. But that is upcoming in the very, very near future.

Any questions around that, before I go on to potential stuff? Okay, not seeing any, a number of us talked. A number of topics have come up over the last – in the past. And so we started to document some of these. And we wanted to run this by the Caucus to get any feedback. And there's been a few people who offered to lead some of these discussions and/or efforts.

---

So I guess we'll start with the first one on here, which is anonymizing query and statistics. Warren, do you want to field any questions or give any thoughts on that?

WARREN KUMARI:

No, because I can't entirely remember what it is. As far as I remember, it was when there are things like [Diddle] runs, how should data be anonymized in the [Diddle] run and in the data output? And should there be a standardized way? Ought it be a standard thing, something like an HMAC or something in the ether, with a secret. But I think that that was it [crosstalk].

BRAD VERD:

Yes, thanks. Paul?

PAUL HOFFMAN:

So, yeah, we keep hearing in various things that we're dealing with [Diddle] data and such. If we anonymize, we would feel more comfortable about publishing a larger set, or working from a larger set. And reproducibility is a real big thing, especially with [Diddle] data, which is hard to reproduce from year-to-year, that we see wild things. If people were willing to work on this, I think that it would, in fact, help the research community a huge amount if we had a standardized way that people agreed on, even if we have to redo it later. But to say, "The way we are doing things in 2017 for anonymizing is such," that would actually help get more research out.

---

BRAD VERD: Great. Thanks, Paul. We will... It sounds like this is something of value that we should look at and come up with a statement of work. So, Paul, can I count you in on helping on the statement of work? Keith, great. Warren?

All right, next topic was terminology. It turns out that we identified – “we” being a number of us – we identified that a lot of us talked about the same thing with different terminology, and that it’s interpreted differently and people walk away with different perceptions. So we thought that maybe we could spend a little bit of time on a document and – I don’t want to say create a glossary of terms, but essentially create a glossary of terms to make sure we’re all talking on the same page, the same wavelength.

Any thoughts or questions there?

UNIDENTIFIED MALE: Sorry to keep getting up at the mic, but five people turned around and looked at me. Given that I’m one of the co-authors on RFC 7719, which is the DNS terminology document, which we’re going to talk about tomorrow in DNSOP, because we are doing a [BIS], for all the people turning and looking at me, we tried to use definitions in that one from RFCs. I assume that Brad is talking about things for which there are not necessarily an RFC-ish definition.

BRAD VERD: Correct.

---

UNIDENTIFIED MALE: It might be a terminology thing within root server operators. For one thing, RSS doesn't mean "root server system" to most of us.

BRAD VERD: Yes.

UNIDENTIFIED MALE: Especially since I was the co-Chair of the [atom] working group, which was what was supposed to have gotten rid of RSS. I keep seeing RSS. I'm like, "No, no, no, we killed that over a decade ago."

However, what I would say on this is if you root server operators would do this on your own – that is, without the help of the outside people saying, "No, you're using that term wrong" – if you could actually come up with a terminology document of how you believe, and it was public, you said, "This is the way we talk," it would be a great thing for us to reference from the 7719 [BIS].

One of the things that we hit in DNSOP is there are a bunch of terms in 7719 which we explicitly punt on. We say, "There is not consensus on this." RFC [foo] says this, but everyone else is saying this, so we don't know what to do. And I assume some of those terms will come up here. It would be great if you could say, "Regardless of what the RFC says, this is what we're using," and we could point to that. Because it is terminology. We're not going to get [crosstalk] –

BRAD VERD: That is the intent [crosstalk] –

UNIDENTIFIED MALE: Yeah. So if you could have something public and if you could finish it within a year, or even less, which is where we're hoping on 7719 [BIS], I would love to reference it.

MATT LARSON: Could you give some example terms for what you're talking about? What would be examples of terms that would be in such a...

BRAD VERD: Go ahead, Lars.

LARS-JOHAN LIMAN: I was just coming to that. The terms that I see in front of me for the RRSAC document is terms that relate to the operations of the root zone, rather than to DNS specifically. That's just along the lines of what Duane said. So how is the root zone for those? What are the barriers, types of operations that are done? How do we reference the various parties that take part in the production? And so on. And hopefully that could mean that we, within RRSAC, could produce documents with consistent terminology. And maybe that could then catch on to other groups, specifically within ICANN, but maybe also within the IETF, if such documents are produced inside the IETF.

---

BRAD VERD: Thanks, Lars. Brian, do you have anything to add here, since you... Brian was one of the key –

BRIAN REID: I'm going to use this mic, because it's taller.

BRAD VERD: Go ahead.

BRIAN REID: I never shut up. As I look around the room, I see probably seven first languages, or eight if you consider Hampshire not to be English. And people do a really good job of expressing what they think they believe in the English that the Internet world takes place in. But at my first RRSAC meeting, there were disagreements that I heard over what "root operator" meant. There were disagreements over what "hosting" meant. If you have a box that participates in Anycasting, what is it? Is it an instance? Is it a mirror? Is it a replicant? Who knows?

And, yes, it is really critical, if we're going to write specifications that last long enough to be useful, then we all have to agree on what the words mean. And writing down what we think they mean is a really good first step. I had actually raised my hand to volunteer for this at the RRSAC meeting, and somebody said, "Oh, you should bring this up at Caucus." So somebody else brought it up. But, yeah, it has to happen. And even if there's only ten words in it, the meaning of those ten words is important. Think about the Orthodox church splitting way from the Roman church in 1100-something over the meaning of two words:

---

Filioque. It doesn't take very many words to have a big disagreement about.

BRAD VERD:

Matt, did that address your question about some examples? Another one would be "root server." I think that means different things to different people. It can be a single server. It can be a letter. It can be... There's many different examples of terminology. And as we went back through some of our discussions and documentation, we noticed even more. So we just felt it was important that we should clarify this.

Paul?

PAUL HOFFMAN:

And non-terms, as well. But it's certainly – like when you say, "What is a root server?" You're not going to come up with one answer. But if you come up with three and you call the first one A, and the second one B, and the third one C, in some document that you're using, where you do it, you can say, "This is root server type A. This is root server type B." The IETF hates doing that, but it's really clear. So feel free to do that. Don't try to adhere to the IETF way.

UNIDENTIFIED SPEAKER:

[inaudible].

PAUL HOFFMAN:

Yes. I said don't use letters [inaudible].

LARS-JOHAN LIMAN: Good point. And I think we should have non-terms as well, because we run into situations where people who are not quite familiar with how it all works try to use terms that they've heard that are not meaningful to people who are into the – like the term “mirror” that I've heard.

UNIDENTIFIED MALE: Clave.

LARS-JOHAN LIMAN: Clave. Yes, there you go.

BRAD VERD: All right. Any other comments or questions around the terminology discussion? I think that one should be fun.

And then another outcome – Lars, I'm going to look to you – is we had a discussion of how things work. This is – well, I'll give it to Lars [inaudible].

LARS-JOHAN LIMAN: Thank you. So as you say, this is another outcome of our discussions, and it's also an outcome of engaging with people, especially within the ICANN circus, where there are a lot of misconceptions about how the DNS work, but the specifics of how the root server system works. And Daniel has produced a couple of documents, published through the Internet Society series.

---

UNIDENTIFIED MALE:           Ages ago.

LARS-JOHAN LIMAN:           Ages ago, which is actually one of the points. They are now 11 and 12 years old, and the Internet has moved on somewhat. Much of the basic stuff is quite correct, but I think it's time to take a new look at them to see if we need to adapt what's in there, if there are new ideas that can be brought into it.

I think we should also try to reach out to other groups, I would say specifically within ICANN, but I am happy to open that up to more arenas, to look for input for how people would like to have that information published. What would be helpful? A public series? A video? Text? Webpages? Animations? What? And we should take a step back and zoom out to try to find ways to convey information about the root server system specifically, but also the DNS in general, how it works, so that we can alleviate some of these misconceptions that are leading people the wrong way.

BRAD VERD:                    Thanks, Lars. Brian, do you want to add something?

BRIAN REID:                   I think a document like that needs to include a specification of why things don't work some other way. My favorite example is I was at a family event, talking about the root system. And my brother said, "Why

---

don't you just use BitTorrent to distribute the root? It's much faster and much more general?" And in fact, I said, "Probably that would be better than what we do now, but we have our traditions." You might want to try writing down why some of these things don't happen.

LARS-JOHAN LIMAN: That's been talked about before, the BitTorrent thing.

BRAD VERD: I think, just for color, for those of you who don't attend ICANN meetings, we, in RRSAC, spend a lot of our time informing the community how things work. And there are some real misconceptions. And when I mean misconceptions, I mean misconceptions on how things work. And we spend a lot of time trying to right that ship. And so this is just the beginning of that effort. And this is how things work as it relates to the root, obviously. Limited to scope.

UNIDENTIFIED MALE: Give an example. Two different ICANN meeting, I have told the ccTLD [inaudible] that it's not really useful to color code top-level domains. And that's just the top of the iceberg.

BRAD VERD: Yes, that's a good example. We've heard like things around the root. So anyways.

---

UNIDENTIFIED MALE: [inaudible] root server.

LARS-JOHAN LIMAN: Yeah, when the root servers stop carrying all the traffic.

BRAD VERD: Right, the roots carry all the traffic. The roots dictate where all the traffic goes. There's a list of myths here that we need to work through and document.

UNIDENTIFIED MALE: [inaudible]

BRAD VERD: That's my favorite, yeah.

UNIDENTIFIED MALE: What we need is we need ICANN Mythbusters.

BRAD VERD: Slow your roll, all right? Okay, thank you, Lars. Lastly, we're going to talk about tools. Wes, do you want to share some thoughts on that, that we've discussed? You did that on purpose.

---

WES HARDAKER:

Yeah. So really quickly, one of the things that came up in our conversations were, what more can the Caucus do? If you guys are bored, we wanted to make sure that you guys had plenty to do. So besides just writing documents, I had this idea of, why can't the caucus do things like produce other useful things besides words on electronic paper? And in particular, we have all this great RRSAC 002 data that most of the roots are now publishing.

Now I'm going to take the microphone out.

And the wonderful things about the 002 data is it's supported by nearly all the roots now. And there are timelines, I'm sure you've seen, for when the rest of them are going to finish it. And it contains significant interesting information on a daily basis on trends, and things like that. So I was wondering, can we produce common useful, helpful tools? Would that be something that the RRSAC Caucus would actually be interested in taking it on?

So just as a reminder, and we can come back to this in a minute if people actually want to do any sort of brainstorming right now, because I think we have the time, but this is the list of metrics that's defined in 002 today, which is the load time, the zone size, the traffic volume, traffics sizes, [R code] volume, and unique sources. So is there interesting things that we can do with that?

And in particular, the things that I came up with were some starting tools might be good. We could start with simply the collection of it across all the root server systems, go get all the files and collect them, and then store them in some way that is useful to people in some sort

---

of database of some kind. Possibly do some aggregation. Possibly do some analysis. And even maybe produce some warnings and monitoring stuff out of that. And then maybe collect a repository of tools so that people don't have to reinvent this every time somebody wants to go do something, RRSAC-related analysis. They can just get the stuff and say, "I'm going to do a new form of analysis. I can run this collection into this database," and then it takes care of half of the work. And then, of course, rinse and repeat with experiences as time goes on.

So my thoughts are – go ahead and go back to the metrics one. Yeah. Is this something that people here would be interested in doing and collaborating together for the RRSAC Caucus? Yes, I see at least a couple of nodding heads. And if so, are there people here that want to take on particular aspects of it? Do we want to create a group? GitHub account? There's a thousand ways to do this, and it's really up to you guys, not me. So I'm standing here, just asking the questions.

JOHN BOND:

Yeah, I think this is useful. We already see a lot of people doing research on the RRSAC data. And so, yeah, it's great to have stuff centralized. Another example would be the way instances or nodes are named is quite different between each root server operator. And I know there's a lot of researchers that normalize that. That would be a good thing to go in this root repository. And as to where, GitHub seems like a logical location to me.

---

UNIDENTIFIED MALE: Just to John's comment, we are already are centralizing all the data. We're scraping all the RRSAC 002 data from operators that are publishing it. So we're compiling a set of that. I think we'd be willing to help support what you're trying to do, which is great.

WES HARDAKER: Is that data currently collected in an OARC-specific manner? Or is it a public repository?

UNIDENTIFIED MALE: It's just a strict mirror of what's published by operators. Only to OARC members at the moment.

WES HARDAKER: Because the rest of the data is public, so whatever we produce here should be public. That was my...

UNIDENTIFIED MALE: Right. I was going to say, after John's comment, I was going to remind that RRSAC 002 data is cluster-wide. There's actually no node-specific information available within it. So the stuff about how operators name their nodes isn't relative to RRSAC.

UNIDENTIFIED SPEAKER: [inaudible]

---

WES HARDAKER: Yeah, to repeat that, there are already multiple people that are doing name aggregation. So why not at least collect what has been done in a common repository kind of way?

Okay, can I see a quick show of hands of who would be interested in contributing to such said code-based collection repository system kind of thing?

UNIDENTIFIED MALE: [inaudible]

WES HARDAKER: Okay. I'm not looking for names. Just there's certainly interest, so we should find... Anybody willing to lead that effort?

BRAD VERD: We can take it back to the committee and find... If you don't want to sign up right now.

WES HARDAKER: I'm more than willing to let somebody else take it over.

BRAD VERD: Yeah. We'll put together a statement of work and come up with goals and whatnot.

- 
- WES HARDAKER: Good. I think that'll be useful. We'll head down that path. Perfect.
- TRIPTI SINHA: Well, continuing on the theme of what more can the Caucus do, as you've seen, there's work that's been done. There's work that's underway, and there's potential work. Is there something the Caucus feels that they would like us to focus on, beginning with us and ending on your plate? But is there something that you would like to send our way and say, "Have you considered a work party for this, that, or the other?"
- BRAD VERD: I heard one earlier that John just mentioned, which was the naming schemes, as far as the site naming scheme. That sounded like one.
- DANIEL KARRENERG: I'm not going to suggest anything, but I'd like to relate general experiences with motivated groups like this. And that is that it's the kind of – at first, it's quite focused and concentrated on a few work items. And then there's a period of enthusiasm. And I think at the beginning of the period of enthusiasm, where all sorts of things – once you meet here or you're outside of your work environment and so on and say, "Oh, yeah, that would be cool to do." And the next phase is frustration, because nothing gets really done well. So what I would encourage or the comment I'd like to make is that we should prioritize. And I think the committee and the Chairs should work towards a way of actually capturing the priorities of the committee and the Caucus, and not take
-

---

on too much at the same time, because I've seen it go bad a couple of times. So prioritize and don't take on too much work at any one time.

BRAD VERD: Thank you.

TRIPTI SINHA: Any other comments, questions, input?

BRAD VERD: Great. So we're done with work then, and work products. We'll go to any other business. Is there any other business that people want to bring up? You're already on the agenda, Daniel. So any other business?

Okay, yeah. So we have a couple action items I guess we'll cover here, which was Tripti and I are going to go back about the Caucus schedule at ICANN and come up with something and send that out for thoughts. There was clarification on the public comments from Jaap that we'll address. And then there's all the other work stuff that we talked through here on the agenda that we'll work in the committee to come up with some statement of works and prioritize them appropriately.

Did I miss anything? That was all I had. All right, with no other business, we will conclude just in time for the social at 5:00. And I think Daniel would like to take over. So thank you all for coming. And the next Caucus meeting that is officially scheduled is the IETF in the spring, which – is that Chicago? I think it's in Chicago. That's the next one. So we'll see you all there. Thank you.

DANIEL KARRENERG: For those people who are interested in souvenirs, they are here. I don't think we're quite meant to drink them here, because otherwise the hotel might charge a [inaudible] fee. So they're really intended as souvenirs. It's Postel, which is an abbey in Northern Belgium. And they used to brew beer there a long time ago. And one of the local breweries, like a decade ago or so, decided to revive this. I have them in three different varieties. One is Blond, which I have the most of, which is the lightest and most similar to a lager. And there's the Double, which is more dark and sweet. And there's a Triple, which is really, as I think it's 11% or something like that. No, it's seven. It's only seven. No, that's the Double.

But it's more like my idea was to bring them to whichever meeting I could reach by car, which this one is, just for people to take home if they are interested. While walking out, a bottle a person.

BRAD VERD: Thank you, Daniel.

**[END OF TRANSCRIPTION]**