
BARCELONA – RSSAC Caucus Work Session 2
Sunday, October 21, 2018 – 10:30 to 12:00 CEST
ICANN63 | Barcelona, Spain

BRAD VERD:

Alright. Welcome back, RSSAC Work Session 2, where we are going to be spending ... Actually, let me just start that. This is a bit of an experiment. We had available time slots and we wanted to try to spend some time with current attendees of ICANN to maybe get some more input on the two existing work parties we had going on and we have the work party shepherds here in the room for the first one and we have one later today this afternoon at 1:30. The goal here was to share what's going on in the work party and try to get some input from anybody here in the room and then that data would be taken by the shepherds back to the work party and shared with them going forward.

So, this might be a very quick meeting, but hopefully we'll get some valuable data and do some sharing. With that, I'm going to turn this over to our work party shepherd for this session. Liman, if you want to.

LARS-JOHAN LIMAN:

Thank you. So, we're going to try to just ventilate what's going on, and hopefully, in the process, get some more input. So, this is regarding the ongoing work party regarding service coverage for the root server system. In a way, this is a restart of a previous effort where we were looking at geo-diversity which didn't really come to conclusion. It kind of grinded to a halt. So, we're trying to look at this from a slightly

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.

different angle and try to put some new vigor into the process. With moderate success, I'll add, but let's at least try.

So, we now have a fairly new work party. We had our first phone call two weeks ago, roughly, and we still haven't found a work party leader for this, so I'm the shepherd from RSSAC, so I'm the [inaudible] for this particular meeting.

Looking at the scope for this work party, you have a link in the Adobe Connect window, the first one. So, the scope is to look for help in guiding principles for root server operators when they deploy Anycast nodes in order to provide good service to the entire network, as far as possible, and to get principles and tools – to solicitate principles and tools for how to measure the service and how to assess whether service is good enough or needs to be improved in various regions, topological regions, on the network.

So, there are actually four bullets to the scope of this working group. I'll read them. They're short, so I'll read them, so you know what the underlying thing here is.

The first one is to explore the concept of accessibility to the root server system, what the indicators, performance factors, and coverage, define adequate and/or inadequate service by the root server system. What factors influence the risk of service degradation or outage, and what are the thresholds that define adequate service. So, the first thing is to identify what can we and should we measure and what are the limits that we need to look at.

The second one is to suggest procedures and/or provide tools that can be used to determine poor service coverage areas or areas susceptible to high risk or lost service.

So, first, what can we measure? Then, how should we go about to measure that? Then, actually do some measurements. So, the third bullet is to identify poor service coverage areas, and given the results from the previous two bullets, help the RSSAC by identifying actual areas where the root server system provides insufficient service.

Now, when we have all the data on the table, the fourth bullet is to recommend to the RSSAC and root server operators how to enhance the service in these areas that have been identified as receiving insufficient service.

We've had the first teleconference. And as I mentioned in the previous session, it wasn't very well-attended. We had five people on the call. We had one person talking in addition to myself. But, we decided to try to collect some data. We've had a little input on the mailing list. So, the second link in Adobe Connect points to a Google Document with some ideas and inputs that have popped up. That was not the document I was intending to look at.

So, some thoughts that have been shared in this context is that while the normal parameters of availability, latency, and accuracy would probably play a major role in whatever we are going to do, but also things like how resolvers choose which specific root name server to prefer and how they [fail over] between different root server instances. There are also thoughts about how to tie this into – how to phrase it?

Into measurements and grading by ... If we were looking for a set of – the term used here is sensors, points on a network where we can perform measurements. We probably need to look for and avoid sensors where there are obvious connectivity issues, which is not necessarily a root server problem. Also, that we could look at the service-level agreements that the gTLD operators need to sign up to and see if that can give any guidance.

There have been pointers to other investigations and other technical papers that relate to Anycast service and there are probably a couple of good ones there that we should pick up and read on and see if they can relate to, give us any good input.

When it comes to suggesting procedures, there's reference to there are tools out there now, if we can combine them and use them in various ways. The notion here that if a certain user or group of users somewhere feel that they don't receive sufficient service that there should be ways for them to easily deploy a sensor with a tool to help understand what the problems are in that region.

My personal thoughts go to the [ATLAS] system which are quite easy to deploy and so on, but again, we need to make sure that the [ATLAS] tests we perform are consistent with what we want to get out of it.

There are references to other various papers here. Of course, the obvious observation that we need to define poor service. In order to hold a situation to a bar, we need to know where the bar is.

Yes. And also to look at what possible effects would there be from other types of root service deployment in addition to the classic root servers that we all know and love. For instance, the idea is proposed in RFC 7706 with operating root servers on your local resolver. I think that's actually a valid comment, to see if there are other ways to mend problems than by just deploying more Anycast service.

Then, this document contains a list of related work.

That's the current status. We are looking for a work party leader. That's possibly going to be a difficult task, given the low number of people showing interest in the work party. But, Ryan suggested that we use this as guinea pig to actually open it up to the full caucus and see if we can get more input that way. I would be [inaudible] to that. But I would like to hear input on that as well.

By that, I open for comments. And let me see. I think that ... Who is behind there? Geoff, please, and then Brad.

GEOFF HUSTON:

I'll admit to not having volunteered on any work caucus anything, so I'll apologize up front. But, I'm wondering if this isn't a little bit of a poster child of our mission sometimes is a hybrid of the political and the engineering and we are in the vast majority engineers. So, unless I'm wrong, this question is political. No engineer is saying, "Hey, latitude 42 longitude 13 by whatever other region is underserved." This is a matter of some residents of parts of the world that feel they are not served as well as others.

So, we are addressing a political question with an engineering answer and I'm wondering whether, one, that doesn't make it less interesting to people. But two, it doesn't sort of lead us in funny directions.

I'm looking at this as somebody who comes more from the economic political science side and my first point would be, well, if you overlaid political borders with the number of instances, you would either settle much of the question where people go, "Oh. I had no idea the coverage was that good there. Thanks." Or, you'd have people saying, "See? That's what I always said. How come there is no coverage in these areas?" Then it would give you something to sort of work on. Brad's hand is up, so I'll yield.

BRAD VERD:

Specifically to your comment there, I think you can interpret this as a political question that we're trying to answer that can be answered with a technical answer. The reason I say that is that we've been asked numerous times. We were asked just yesterday in the how it works, "How do we add a server? How do I get a server?" And the question I handed back to them was, "Do you need one?"

So, back to what Liman said, it would be nice if there was a test that you could provide some of these people to show the availability, the service ability, the reliability, whatever word you want to use, of the root system where they are.

Also, the questions we got yesterday in the How It Works, it would be nice if the test could show – I'm trying to think of how to word this. But,

most of the questions were these are out of the root server system's control. This is beyond our ... We can't control the routing system. We can't control your resolvers. We can't do this.

If you look at the gTLD monitors, the system that ICANN put up to do their SLAs – and I'm just going off of my experience on dealing with the contract side of that – they have sensors all over the place. They probe the gTLDs directly and that is the measure of availability for the gTLDs. It's not through the resolvers. It's directly to the servers, which I would imagine would be the way you would – the only way you could really give it to somebody as their own little test because that's really what the question always comes down to is "I want a server" and if you give them a test ... It would be great if we could give them a test, I think, to say, "Here's what service looks like from where you are." That doesn't exist today, short of somebody with some DNS knowledge sitting down and doing it.

Clearly, based upon the line of questioning that we get at every How It Works meetings and caucus meeting and open RSSAC meeting is there's not a clear understanding of where the responsibilities are. Everybody thinks the root. Everybody at the root – that we have all the control. And we don't. So, it would be nice if this test would help with some of that. I don't know if that's possible, but it would be nice.

I think it would be ... In an effort to not recreate the wheel, I hope we're looking at the gTLD tests. I hope this group does. I think that's a great place to start because ICANN has spent a lot of time and effort putting

that together with equations and SLAs. And I'm not proposing SLAs. I'm just saying that there's a lot of effort put in there.

And going to what Geoff said yesterday, it would be nice to just put a line in the sand. Not a line in the sand as to what this is what looks good, but just here's a list of tests that could be run from your location, and start reporting on what those numbers are, if that makes sense. You don't know what you don't know, right? It might be really good. It might be really poor. But, the question of – the definition of what is good and poor depends on I think your location and where you're asking that question from. I don't know if we'll be able to define good and poor, but I think it would be really good if we could define a set of tests. Under 100 milliseconds might be really good for me, but if I'm doing – I don't know. If I have some weird requirement or know somebody who does, maybe they want some 20 milliseconds or something. If that's the case, then you either need to move or we need to get more [service] in there. Do you see what I'm getting at?

GEOFF HUSTON: I'm understanding increasingly why we need to have a diversity of approach.

BRAD VERD: Yeah.

GEOFF HUSTON: Because, no, frankly. There's a boiling the ocean thing going on here where it can't boil the entire ocean at even rate. It's not worth looking at the wave that's about to take your house out. So, there's a part of me that just says it would be interesting to take an allegedly underserved portion of the world and just look at it and say, "What is the latency here? What are the receive times? Are those reasonable given the infrastructure in the area?" And I think that would be a better data point to somebody who comes in and says, "Why can't I have a letter?" And we can say because we believe—

BRAD VERD: That was just an example. I wasn't trying to ...

GEOFF HUSTON: No, but I'm just saying. We keep answering political questions at those forums with technical answers, and if you ask the people asking those questions whether they were happy, I guarantee none of them have ever been. So ...

LARS-JOHAN LIMAN: I would like to toss in a small thing here. I can see this as a way to turn this question, which I totally agree is political and technical, both. But, if we can pull this to some kind of conclusion, we could diffuse it partly by saying this is not a technical problem. If we, by these measurements, can find that we have very few underserved areas or poorly served areas from a technical standpoint, because this is obviously a political problem ... If you want, this typical question of, "Why can't I have my

own server?” We can say if you want that, you’re not trying to solve a technical problem because we believe that we have that in control. So, you are addressing this from a political standpoint. Right. Then you can shape the dialogue into a political one which is a very different beast and needs to be treated differently but at least then we have diffused the technical part of it.

So, I believe that Brad was first and you second, but I’m not quite sure. Okay, Geoff?

GEOFF HUSTON:

To whom would you like them to bring the political questions, if not RSSAC?

BRAD VERD:

If I may, and I think I’m going to reword maybe what you said, which is I think it was ... So, yesterday, we were asked about why ... How do I add a letter? Then, Fred turned around and said, “With the 1,000-plus servers we have out there, do you need one?” It would be nice, in my opinion, with that political question that was just asked to turn around and then askk, provide empirical data that says, “You’re not underserved.” Right now, we didn’t. We said, “Why?” This could turn around and say, “Here’s the empirical data that shows you’re served very well.” Not saying it’s ... Yes, it’s answering a political question with a technical answer and I know there’s a lot of passion on either side about, “Should we do that or should we not do that?” but that’s the world we’ve lived in here now for a long, long time. So, yeah, go ahead.

GEOFF HUSTON:

Look, this area of empirical data and what you're trying to measure is incredibly complex because most users do not directly query any root server. They go through the recursive infrastructure, and if your recursive infrastructure includes folk doing aggressive [inaudible] and things like that – and here I'm referring to all of Africa and all of Google, they actually get surprisingly good service. Surprisingly good, irrespective of the density of root servers located close to them. So, ICANN's experiment that directly queries against gTLD service would show a shocking result, but the users see a really good result.

So, part of this trying to understand what's the service and who are you trying to serve is also trying to understand who and why do folk use directly the root servers and what's going on? We have a suspicion I think around this room that most of us spend most of our time shoveling garbage. We actually don't understand there's any users behind all these queries. We just don't understand. And it becomes very difficult, Brad, to sort of put down a solid technical case around metrics and so on when we actually don't understand what part of the current metrics is real and what isn't. So, I'd be all for doing user-based experiments and it would oddly enough lead to a very different kind of answer.

There is a second dimension. Another letter would solve it. I don't want to go there. I think it's an entirely different discussion, Geoff, with all due respect, to simply look more at, "Is the DNS service infrastructure supplying root domain answers efficiently?" is actually a more subtly

focused question than, “Are there enough Anycast root servers out there?” which kind of assumes an answer that you need one, whereas the reality is most of the time, most folks don’t.

And if you looked at root service delivery as an artifact of the DNS, you might be on more solid grounds to at least make a technical case. I agree, the political thing will never go away. All the technical thing will do is boot the argument that we get shocking service because ... Boots that out of the door because you don’t. You get really good service, your complaint is of a different dimension.

BRAD VERD:

Yes, I agree. That’s all I’m trying to help with, not even accomplish. I don’t think there’s a perfect answer, much like you said. Just like we had yesterday the question, “Why did I have an outage?” It’s like, well, there’s a million different reasons. There’s so many different variables, so many different things, it’s not a black and white type of question you just asked, but they believe it is. Most people believe it is. I shouldn’t say “they”. It’s complicated.

So, it would be ... I don’t know. Maybe I’m in my own Narnia here. It would just be nice if there was something that you could hand them and say, “Run this,” and maybe there’s a series of test that run directly against the root, so you can show them, “Here you are in relation directly to the roots. Now let’s do some queries using your resolver.”

What I don’t know – and I was going to ask this question. I don’t know – how do you do a [KS test] query to a 7706 root and figure out where it

is? I don't know if you can, right? So, that presents a bit of confusion because you might not ... I mean, I might be [inaudible] be anywhere versus I know all the root support to [KS test] query, so you could figure out if somebody is in South Africa and they're getting a root server that's in Europe, okay, that is probably an ISP issue that you've got going on there.

Ultimately, I think it would be nice, if that's the right word. It would show us trying to help the user, help educate the user, but we're certainly not going to be answering the perfect question to say ... I think the answer to number four in your presentation is we're going to find nowhere is going to be underserved, unless somebody is like, "Well, Antarctica is underserved." It's like, okay, you win. You got me. I think that's where we're going to end up. To me, that's like where success on this experiment, so to speak.

I don't have the empirical data to support that. Perfect world, we're just not going to get there. It's just trying to move the needle on some of this stuff and/or provide defensible arguments like yesterday where we were like, "Do we need one?" My answer – and it still holds true today when people say, "Can we add one?" Yeah, but we could cut it in half, too. Eyes kind of blow open and people are like, "What? What do you mean?"

GEOFF HUSTON:

This is a little frustrating because what we're sitting around here saying is the current root server system is pretty [inaudible] adequate and nobody around the globe is underserved. I buy into that.

BRAD VERD: Much like people say we need engagement with the caucus, but people keep saying, “I need one, or I’m underserved,” so it would be nice to have.

GEOFF HUSTON: There was a phrase in the response from the board that was something like, “Obviously the system has failed.” I’m getting it wrong, but the implication was that this is falling apart. There’s a real disconnect there. What the threats are to it functioning are not that it is currently not functional, although we’re shoveling nonsense as Geoff points out. We’re good at it and apparently this is a signal to noise ratio that works for something. The existential risk to the system – sustained [inaudible] attacks and all those ... Those things are real, but I just wonder how much this is a matter of not understanding something that’s complicated and I think the questions tend to be political and our answers tend to be technical and that’s part of the whole Mars-Venus disconnect.

BRAD VERD: I think, using your body language there, I’m not trying to do this. I’m just trying to do *this*. I don’t think we can connect with a technical or a political question. It’s just not going to happen. But, if we bring it closer together, it makes it easier for us to give that answer and explain it.

GEOFF HUSTON: I'll leave it with a caveat. I think if we think a little more about putting a little more of a political answer in occasionally, it wouldn't hurt, because otherwise, we're just Root Ops.

BRAD VERD: If it's implied that I said we shouldn't do that, I'm not trying to do that at all.

LARS-JOHAN LIMAN: Because that's what we're really good at. [inaudible], I mean.

BRAD VERD: Going back to what I said earlier, it would be nice if this tool did some sort of – or could do, let's just say some sort of chaos text query so you could figure out the physical location of where you are type of thing. Don't know if that's possible. I'm kind of picturing DNS – some DNS tool that somebody could bring up on their computer and spot out a report type of thing.

UNIDENTIFIED MALE: So, maybe – [RIPE ATLAS] gets you pretty far in this way. Maybe the work party could explore the idea that if somebody comes to RSSAC and says, "Hey, I'm underserved," RSSAC says, "Here's a [RIPE ATLAS] probe. Put it on your network for three months and we'll evaluate it. We'll see what the measurements say."

BRAD VERD: Yeah. I think trying to find a partnership with something like that, that's one approach and it promotes the ATLAS monitors which are used by a lot of people which I think is good. So yeah, I think there's a number of different things that could be talked about and hopefully the work party will do that.

LARS-JOHAN LIMAN: I also think that the [ATLAS] people would be open to ... If we come up with an adjustment of a test, a profile if you wish, that would further help us to get this information that we see that we need to judge this or assess this. I think the [ATLAS] people would be quite open to [deliver] on that. As Duane says, it's probably mostly already there.

UNIDENTIFIED MALE: As someone who specializes in end user tests – and I do around 10 million a day – [ATLAS] seems woefully inadequate. There's nothing in Libya, nothing in Niger, nothing in Chad. There's one in Mongolia. The folk who are complaining are not the folk who have [inaudible] infrastructure and [ATLAS] really isn't your answer. It is possible to reach users in all of those places quite legitimately. What is not possible to do is actually see the result of the root server.

So, it's incredibly easy to see the user with a domain that does not exist and get them to ask it to get that very same [NX] domain response that you're so good at shoveling.

And if you could see it at the root – any root server, you'd be a lot better off in trying to answer, "It went this fast. It came through your

infrastructure this quickly.” So, part of it is there are a lot of tools at your collective disposal that are way much bigger than [ATLAS] and way more effective at doing this, but it requires both ends. It’s not just one end probing into the middle. If you want to make this work, you’ve actually got to see it at the other end as well and connect the dots.

BRAD VERD:

Well, I hate trying to ... We shouldn’t try to engineer on the spot which is what this group does because we’re a bunch of engineers. But if there was ... If I’m a user in Africa somewhere – I forgot the name of the country you used, but ... Chad, thank you – and I had a webpage that I went to, you’ve now created that connection. If that person initiates that test, then you could come up with those answers that you – and those answers could be reported on the spot to the user and collected on the backend and we could start putting that together. It would be nice to have a tool like that, I think. Maybe I’m wrong. Maybe I’m naïve, but that’s just ...

And it seems like most of the ... It seems like we’ve done most of the work already, just haven’t maybe integrated it into a test for the end user to do. Geoff?

GEOFF HUSTON:

So, obviously, if we had an app that was [HowsMyRootServer.com] or whatever that somebody could run from anywhere and when they say, “Why don’t I have one?” or “I need one” and we say we can run it, is that what you’re talking about?

UNIDENTIFIED MALE: [Is that crazy?]

GEOFF HUSTON: No. I'm just trying to think where we come up with a couple thousand bucks it would take to pay somebody to integrate the existing tools into enough of a packet. We're not talking about my grandmother using this. We're talking about technically savvy people. It shouldn't take that much.

UNIDENTIFIED MALE: I haven't seen my hand. Sorry. Not high enough. Then we'll come to [inaudible].

UNIDENTIFIED MALE: I've been spreading around [inaudible] [probes] at ICANN meetings for a couple of years and especially to the regions you were talking about, Africa and Mongolia and stuff like that. All these things seem to end up in a black hole. You never see them back. So, the [inaudible] to do something in this region [inaudible]. I don't have them with me anymore because [inaudible] at the [border].

One thing that might help is that the [ATLAS] folks are now working on having [inaudible] machines doing [S probes] so that might be easy to spread around than [inaudible] people can use [other stuff]. So, that might help a little bit for the [inaudible] stuff.

Then, there is this project of [inaudible] which actually looks now for more than a year to various aspects of, especially [resolvers] but it's probably easy to expand it to specifically look at [inaudible] as well. [inaudible] contact with those people and I already spoke with them and said [inaudible] stuff in it.

But yes, it is difficult to find, to get [inaudible] started, dedicated [inaudible] Geoff is doing. And other sources we found out that the [inaudible] to find stuff is using – there's some open source VPN service and we have been using that as well and that's spread out way more places because it's the end user starting to [inaudible] to various [blocks]. So, there are some ways of doing some measurements, but they are limited. But there are way much better [inaudible] than [ATLAS] because that's naturally bias. We do our best.

GEOFF HUSTON:

So, I like the idea of giving the people that really are underserved the ability to prove it and come out of the closet wherever you might be and really show that need. The flip side of that is, again, if it's really a political problem, then I'm going to take this [ATLAS] probe, throw it 1500 hops back in my network and show you really poor data. So, how do we balance that accuracy [inaudible] measurements?

BRAD VERD:

I think the [ATLAS] probes are one source of data. I think if you give somebody a tool that is kicking off these queries directly from the

machine that they're on type of thing, you're no longer using [ATLAS] probes.

UNIDENTIFIED MALE:

Right. So, it's more of where do they deploy that sensor, if it's [ATLAS] or any on their laptop. I mean, if it's over a [RS232] connection, they're going to get poor performance. One of the things that was mentioned was the [verve plotter], the tool that can go out and ... That's an alternative to maybe [per/24 or per/AS], do measurements from the other side, so that maybe we're not reaching all the way back to the infrastructure. We just need a way to separate the root server problems from end user [inaudible] problems.

GEOFF HUSTON:

Look, the real issue is actually measuring the user's recursive resolver infrastructure. How fast does it take to get a no? How fast does it take to get an answer? Whether or not [inaudible] into the root servers or not, in the grand scheme of things, the user is completely unaware. It's basically the time from when the query was launched indeed when they clicked until the time something happens.

It's actually weirdly perfectly possible to use a [inaudible] measurement system to measure the [NX] domain delivery rate. Actually, it is indeed quite possible to see the user's resolution system with names that do not exist – have never existed – will cause [a cache miss] all the way through into a root system and time how long it takes to resolve the fact that it's a no.

Now, if it's aggressive [inaudible], it will be really quick. If it goes back the root, it will be slower. If you think that measurement would help you all over the world for every AS, we can do this. I can do this, if you think it will help. If you think it's a waste of time, I won't bother because it's a waste of time. But, I'd really appreciate the feedback. If you think this is promising, we can head down that direction and construct a massive long-term experiment that actually weeds out [NX] domain delivery response. An [NX] domain that realistically refers to names that just aren't in the root. So, is it helpful?

WES HARDAKER:

So, that's for mentioning [inaudible] a second ago because I've been trying to raise my hand for a while to talk about it. We developed that ... No, it's okay. I was being very shy about it. It's not your fault, it's mine.

Anyway, it's designed to measure response times. We generally get 50% of the [24s]. That's a huge number. I'd love to work with an Anycast system that had a huge number of nodes so that we could actually get a fairly accurate latency map of pretty much the whole planet for a widely-distributed Anycast system, like L or F, for example. It takes some setup and it takes some work, but if anybody wants to work with me, I'd be more than happy to help make those maps and prettiness so that we'd have some real data.

But, more importantly, I think half the goal of this work is not just to measure one system, right? What we're missing is not just the technical answer for, "Can you get to this address really quickly?" It's what do all

the resolvers do? And Geoff is talking about this, too, although he's talking about it from the web point of view and I'd like to see even a greater coverage. Fortunately, web these days is – most of it – is not mail, however.

We need to be able to study that resolver selection issue, too, of how many are they preferring and there has been recent studies in that area, too. But, in order to give a technical answer to a political or a technical question, we have to evaluate the system and I think that's where we're missing the most data.

BRAD VERD:

A couple of things. One, regarding [inaudible], if we could do a latency map for the root server system, that would be phenomenal. I don't know if that's in scope here or if people have talked about it or whatnot. I've said this before. We've spent a lot of time at my company figuring out latency and availability for our TLDs and it was a huge amount of work. I talked about this. I think we even put some analytics in there and tried to estimate how physically far people are from certain servers type of thing.

I think there's a couple of different things that would be nice-to-haves. This is just in my head. I'm just not filtered, so please don't judge too much. That is it would be great to have some global latency map of the root server system. It would be phenomenal if we could provide that or create that.

The second piece is you step down to the user base and a user-base test, to me there's ... And this isn't necessarily trying answer – I mean, it is. It's providing empirical data to the age-old question of, "Why can't I?" "Am I underserved?" I guess. To me, there's two things that you've got to overlay on top of each other.

There's the response time to the roots and what root am I talking to, but then there's also this – and you've got to overlay that on top of the resolver piece because, look, if I do a query directly to the root, here's your response time, if you then do a query through your resolver and the response time is much larger and you're going to a server up in Europe, then that's like, "Here's your issue. You need to go talk to your provider." You're not underserved. You're talking to a provider who isn't configured or has their networked [paired] somewhere that's just weird. But that's not the result of the root server system underserving them. That is the result of their network configuration.

I'm not trying to solve that problem. I'm just trying to ... It would be nice if you could easily bring attention to it. Does that make sense? I see Liman taking notes pretty heavily. Any other discussion? Like I said, this is a bit of an experiment to see if we could get input on the topic. Go ahead, Liman.

LARS-JOHAN LIMAN:

Process-related question. Are there any opinions about the idea to make this a guinea pig and toss it to the entire caucus? Should we? Is anyone opposed to?

UNIDENTIFIED MALE: I'm in favor of. I think it makes a lot of sense. It was a generic topic that was under ... It had too few people contributing and then we brought it up and it was like I can't believe that's not fascinating, at least at some level, to everybody.

GEOFF HUSTON: Now that I see this on the screen right in front of me, I'd just like to raise a subtle distinction of what I think would be useful and what you think you're doing. When you say root server system, I read root service system. In some ways, what the users are doing is not sending their queries to A, B, C, D root server dot-net. They never even know it and they never answer anyway, except [NX] domain. They're actually asking questions of the root service. And even it requires an introductory text about the distinction and the scope of what can and should be measured, I think it's a very important distinction because we're seeing this concerted effort to enlist the entire recursive resolver set into providing answers about the root zone and whether it's just aggressive [inaudible] caching and zone learning or all the way through 7706, whatever you want to call that. The users benefit.

And if we're talking about the roots as seen by the users in resolving names, I see this as a root service problem. I don't think the root servers have to be the be-all and end-all of any answer here. I think service is a better way of capturing there are many productive solutions to this and that may be a helpful way also of phrasing this. Thanks.

UNIDENTIFIED MALE:

So, I heard a lot of really great ideas and I'm wondering if – and this may elongate this work party. And being that it's a guinea pig [as] throwing to caucus versus having just specific members, restarting this entire work party anew with a new scope and then a new ... Of course, opening up to a new group.

So, first, using the caucus to develop the scope and then ... Because as Geoff pointed out some really good, and Brad and Geoff, and then developing the new wording, new scope, and everything and then of course shifting it over to the full caucus for input. Would anybody have any objection to that, as to developing kind of a new scope, letting the caucus have input into the wording of the scope? And yes, it may take a couple months to get the scope correct, but then we could start the work party. And this way, it would kind of ... Because we relied upon the RSSAC members to go ahead and create the scope, which I was one of them, and hearing what Geoff Huston has to say, I realize that maybe we got the scope not. And I'm not objected to modifying the scope and letting the caucus have input into modifying the scope. Anyway, that's about it.

BRAD VERD:

I'm hesitant to restart anything. I'm hesitant. Not opposed to it. I'm certainly ... If that's the right thing to do, let's do it. If we want to change scope, to me the statement of work is nothing more than a starting point, right? So, if scope needs to change, why can't the work party do that and just inform the RSSAC that, look, we've done our initial

evaluation and we feel that maybe this is where we should go. To me, that's a lot easier than let's stop, let's restart. There's a lot of discussion just in the SO/AC leadership meeting with the ICANN executives [or the] idea of in a goal to be completely transparent and completely accountable, we've become completely ineffective. I really want to try to avoid that. I don't want the process to be like, "We've got to change the scope and we've got to follow procedures, so we've got to kill the whole thing and start over." I'd really like to be as fluid as possible so that the work party can get work done. That's my two cents on that.

LARS-JOHAN LIMAN: Yeah. I think there's probably [inaudible] on that. I'm in favor of your way to address that.

BRAD VERD: Anything else or shall we ... Andrew?

ANDREW MCCONACHIE: So, just on a process question, if I understand correctly, we are opening up this work party to the entire caucus list? That's the idea? Okay. So, there's a couple of maybe minutia things that we need to work out together and we can just work on that. Sure.

LARS-JOHAN LIMAN: Yes. Right.

BRAD VERD: And if the scope needs to change, then I think it's just suggesting what that scope is and saying, "We've talked. We've worked on this. This is what we think we're going to go answer and come back." Because what I'd like to avoid is we've had work parties in the past that have gone off, done their thing, came back but didn't answer the question. And didn't change the scope. So, that's a complete disconnect there. So, I'd like to try to avoid that.

If there's nothing else, we will call this early, give you guys some time back. We're here again at 1:30 to talk about resolvers and see ... Again, it's going to be the same type of dialogue, a bit of an experiment, and collect that information and share it back with the work party leader.

LARS-JOHAN LIMAN: Not here, I believe. Not this room.

BRAD VERD: Oh, do we move? Does this meeting at 1:30 move? I'm sorry, we're across the way at 119. Other than that, have a wonderful lunch. We're adjourned.

[END OF TRANSCRIPTION]