
PANAMA – RSSAC Work Session (3 of 5)
Wednesday, June 27, 2018 – 10:30 to 12:00 EST
ICANN62 | Panama City, Panama

TRIPTI SINHA: Why don't we just get started? First, Steve asked for a few minutes, so I'm going to turn it over to Steve. And somebody has got me on speaker. Alright.

[STEVE SHENG]: I just want to say words of appreciation and thanks. My apologies, I had to run early yesterday because I had a family emergency. I think, personally, I'm very humbled by the kind words and appreciation. I think we are very proud to support your work.

I was talking to Andrew the other day and the work that he did really completed what Postel [left behind] in 1998. For us to be able to support that work and see that Postel's legacy and the work completes, it really means personally a lot for me. So, thank you. I think I speak on behalf of the staff supporting the RSSAC, and you will continue to have our utmost support and appreciation. Thank you.

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.

TRIPTI SINHA: Thank you, Steve, very much. My apologies for missing the first session. I don't know if you guys know, there's a Swedish professor who is doing a survey. Have you heard of him? He's going to administer I think Barcelona onwards. He's put together a survey and he's fine-tuning his questions and I'm one of his guinea pigs. So, helping him fine-tune his survey.

This session is about two work items that bubbled up to the top. Who were the two leaders that we assigned? You and?

LARS LIMAN: I don't know who is. I am for the geo thing. I don't know who is for the other thing.

TRIPTI SINHA: The other one. Did you want to take the lead?

LARS LIMAN: I can do that. We are ... How should I kick this off? We decided to go for a work item regarding geographic diversity of root name servers as a spinoff or a continuation of the work with Anycast stuff that happened earlier.

We are currently at the point that we have a draft statement of work. It's not finished yet. It's not something that ... It needs

more work, not just the final touches. But, let me see if I can share. Okay, good.

I guess what I would like to do here is to get some input on you on what things we are missing or possibly should remove of the statement of work.

The original text, which is still included in the document, and just for reference, so it will be removed before we finally use this document.

It says there are five areas to identify, underserved areas. Sorry. That's a bulleted list which has been copied and pasted. It's not in a concise date. It also involved defining what underserved means, and we have rephrased. We avoid using the term underserved in the document as it reads further on. But does not receive sufficient service. So, I'll define what that is. We would like the group to come up with the areas that are not served well.

UNIDENTIFIED MALE: Excuse me, Liman. One second. That's the wrong document. That's mine. So, that link doesn't match. Sorry.

LARS LIMAN: Okay. So, am I on the wrong document or are you in the wrong document?

UNIDENTIFIED MALE: No, you're on the right document. We're all staring at the wrong thing.

LARS LIMAN: Okay. So, I should again try to share this. Give you a link so you can read it. Let me see here. This is a good one. Copy. Let me try and get this one into the chat window. This is the doc I am reading from. I'll give you a minute to find the link in Adobe Connect. Interesting. Sorry, I didn't read properly. I just want to get a shareable link and copy the link. Here we go. Thank you Google docs. An internal error has occurred and your request was not complete. Try again. There we go. Let me see if I need to do a link now. Does it work now? Okay, good. I created this inside NetNode's internal Google doc structure. I shared it with a few individuals and that works, but I needed to do some extra wand waving here to make it available for all of you. So, if you now have access to the document.

The scope was listed as a few bullets or a few statements. We have rephrased those into five more structured bullets. That's where we would like to have more input. Both of you think the

text needs phrasing or if you have more bullets or if you want us to remove stuff here. We will ask the work party to define what's meant by well-served geographic area and also to look at what indicators of performance factors determine how well a geographic area is served. That is an interesting question because typically service is not provided to a geographic in my mind. It's provided to a topology.

So, you can very well have a situation where one ISP received very good service and the neighboring ISP covering the same geographic area has very poor service, depending on how the structures [inaudible] and so on are set up. But I guess that's part of the work party work to look at that and try to define and come up with indicators if there are any such to find. Also, as [inaudible] to explore the concept of how well-served a geographic area is.

Then, to the more sensitive details or interesting parts, to identify areas which receive insufficient service. Now that we, in bullets one and two, have figured out how to measure it. Then as a last [inaudible] actions to reduce areas of insufficient service.

What else do we need to look at here or talk about? Wes?

WES HARDAKER: Tangent. In the past, we talked about having a caucus group, develop a tool that could actually do this measurement, and maybe it would be good to roll that purpose into this one so that we could say it's part of the output. It wouldn't just be a document. But, methodologies and potentially software for performing analysis in [inaudible] regions.

LARS LIMAN: Anyone else? Tripti?

TRIPTI SINHA: I'm just picking up on a point you just made, that geography doesn't necessarily map to topology. That's a really key point here. Do you think the title should be modified to speak to what we are actually trying to do here? We're trying to map the diversity of the service to probably the density of a user base, a community base?

LARS LIMAN: I'm not quite sure. That's why this area needs to be explored.

TRIPTI SINHA: Right.

LARS LIMAN: So, let me take a very simplified example. If there is a nation ... I'll use words here and these are not definitions of ... This is just a thought example here. If you have a nation with two ISPs and one ISP receives very good service for root service and the other one doesn't, and they both cover the entire country and they have half the population as customers each, is that an underserved area? Because the area is well-served to 50%. Maybe that's the approach we need to have. How large fraction of the population is receiving good service or bad service?

It won't come to a case where you can draw borders saying this area is well-served, but it will turn into different shades of gray how well an area is served in numbers of end customers.

TRIPTI SINHA: Yeah. You almost want to come from the perspective of topology. Where's the top logical need for more service?

LARS LIMAN: Yeah. I'm happy to change the title. I kind of inherited the title from previous work. I always try to see it from the topological view. I would be happy to change the title, actually.

TRIPTI SINHA: As Wes just said, even that is a moving target, because as the topology is constantly getting enriched and enhanced, and let's say the other ISP is now resourcing better, then there you go. That changes the landscape again.

LARS LIMAN: But, there are probably also areas that, for a long time, have received insufficient service where things are not changing rapidly and there is lots of room for improvement for other than just fine-tuning.

WES HARDAKER: I think there's always room for improvement. I guess my question, and this is a bit rhetorical so I'm not trying to be snipping or anything. What is good enough?

What I mean by that is I'm going to go way back in time to a bunch of people who were here. Some of you guys were not. No, not in my day. But, we did a presentation. We being Verisign did a presentation of Root Ops which kind of touched on this stuff. It was in the Prague meeting I think maybe four or five years ago. It was a huge effort. It was a huge effort on our research team that studied the resolvers talking to our roots and where they were and where they were coming from and who was talking to us and how often they talked to us.

You could see very interesting trends. There was a huge concentration of resolver queries coming from South Dakota or something weird. Then you did a little bit more research and you find out that Google has a huge data center there and that's why that was coming from there. Then, we did some extrapolation that basically said, well, based upon all the queries we've seen, it's clearly not accurate because it's network topology versus geographic location, but we're going to – we know that they don't overlay, but we're going to assume a couple of things. Then it basically said that on average we believe that is was – every user was at a minimum of ... I forget what it was. It was like 100 miles from ... No, it wasn't even that. It was I want to say 500 feet or something from a root server, based upon the queries we're actually seeing.

So, I think it's going to be really interesting. We need to define what is good enough because you're going to find – I mean, it's easy to look at a map and point to different regions on the planet and say, "Well, clearly, they're underserved." And I'm like, "Underserved as defined by what?" I think that's a question that needs to be answered first before you can say a region is underserved.

LARS LIMAN: Yeah. It is question number two in the document, I would argue. Explore the concept of how well-served are geographic areas. Number one, sorry, define what is meant by well-served geographic area. That ties into you can rephrase it into good enough, but I see that as the issue that you described. And you are quite right about it, definitely. You can measure it in return time. You can also measure it in resiliency and you can measure it in geographic areas with political control and what have you. There is an infinite number.

UNIDENTIFIED MALE: That's [inaudible], which is well-served has multiple angles. I was thinking availability, I don't think we've talked about.

LARS LIMAN: I don't see anymore hands here. I intend to continue the work on this document. Looking at all of you, both those who helped me write them and the rest of you, shall we have it as a goal to have this ready for the July meeting or is that too soon?

TRIPTI SINHA: How much more work do you believe you have in tightening this up?

LARS LIMAN: To be frank, not all that much. It's more than just putting the comments right, but not all that much. I'll be working for two more weeks and then there's the IETF ... You have two more weeks and that gets them to the [inaudible].

TRIPTI SINHA: It would be nice to have it done by July. Would you have to engineer it backwards? Because if we approve it, we need a [one-week stable] document.

LARS LIMAN: Ah correct. Okay. So, we need a [stable] document by the 3rd, if I remember right. Okay. Interesting.

TRIPTI SINHA: Probably want to send out a week prior to July the 3rd some time for people to comment on it, right? Do a back and forth.

LARS LIMAN: Okay.

TRIPTI SINHA: It's a tight timeline. But, if you could shoot for it to be ready by July 3rd, it would be nice.

LARS LIMAN: Okay. Oh, there's a comic strip about that. I won't do it here. Alright. So, we'll try and see how far we get. Steve?

[STEVE SHENG]: Thanks, Liman. Given our experience with the caucus on the last work party on Anycast, I have two questions. From the RSSAC perspective, what do you envision the expertise from the caucus to complete this work and whether the caucus has sufficient expertise? Second, how do we structure this work so that we have a better chance of success this time?

LARS LIMAN: To your first point, what type of expertise we're requesting from the caucus this time. Okay. I actually see two different types here, dealing with different parts of this. One is the technical side, trying to define if there is a technical service level, what is sufficient service, to try to define that. The other one is, given that we have this and are somehow able to measure against this threshold, people who know various areas and can help us pinpoint topologically where the problems are. That's a totally different set of people because that's probably ISPs and operators. Maybe we don't find them in the caucus, and if we don't, then I assume we will receive that feedback and then we will probably have to look elsewhere for that. Or we're just receiving the information that we don't know [inaudible] these

people or those groups could be very valuable. Who was it? Fred? Wes? Duane, sorry.

DUANE WESSELS: Did we really settle the issue of geographic in this work party? I feel like that's something that the work party could get really hung up on.

LARS LIMAN: As I said, I'm happy to change it to something else. I inherited the phrase.

DUANE WESSELS: I could envision that either – well, it could stay as is and focus on geography only. It could expand to focus on geography and topology or it could focus only on [inaudible].

LARS LIMAN: Or, we could say just areas. That's dangerous as well.

DUANE WESSELS: That was my second point, that if you leave it vague, then it's dangerous. It potentially expands the scope of work quite a bit.

WES HARDAKER: I think what we're really talking about is not necessarily geography or topography, but rather, are we providing adequate service to Internet users at large or their resolvers? Maybe we ought to be more generic and say is the service coverage adequate and equal?

DUANE WESSELS: Well, I feel like one of the reasons that we're having this document is because on more than one occasion, like in the information sessions, people have stood up and said, "How do I get a root server in my country?" We need to decide how we want to answer that question? Is the answer you don't necessarily need one in your country because XYZ, or if that's a valid question, then we need to decide what are the parameters for when they get located?

LARS LIMAN: Fred?

FRED BAKER: Yeah. This one actually is Fred. Looking at the question that Duane just asked, and we got that as recently as last IETF meeting – somebody from Zimbabwe wanted to have a root server in their country. The simple answer that I gave them was here's how you would ask for a root server. People back at the

ranch are going to look at you and figure out by what metric they would then go put a root server there. For the record, we haven't done so yet. I think the need for a server is something that has to be demonstrated.

But, from the perspective of an RSSAC caucus activity, seems like the question that you want to identify is are there zones where Internet – DNS access – is not wonderful? You're going to need some kind of a definition of a zone. That might be where are the major airports in the world? I don't know. But, where would you find an Internet user to ask the question about? You're going to have to define that in some way. Maybe that's part of the work statement. I don't know.

LARS LIMAN: So, defining what the corresponding endpoint of the traffic is and how to reach them or how to measure from their point.

FRED BAKER: Yeah, something like that.

LARS LIMAN: Okay, thanks. I'll try to fit that in somehow.

TERRY MANDERSON: I don't want to do the work party's work here. I don't want to do the work party's work here. It sounds like we're starting to get into that frame. Whether we talk about topology or catchment areas or zones or countries or airports, I think that's besides the point. I think one of the things that needs to be assessed is simply an underserved point.

Now, whether that underserved is they really don't have root service or their root service is something different compared to the 95th percentile of all catchment areas, I don't know. But, I would urge us not to get caught up in those discussions here, and actually give the caucus work party the option to investigate that. Yes, it sounds dangerous, but that's why we actually do have a shepherd to help them along. And if they start to rathole down one particular direction, go, "Okay. We have established that this is exceedingly hard to describe. What's the third option? Is there another way to think about this? Can we look at this from a different perspective using a different box?"

WES HARDAKER: So, if I may translate that into a bullet number one, step one should be to define a metric by which to measure underservice.

TERRY MANDERSON \: Yes.

WES HARDAKER: I debated that and I ended up with underservement for two reasons. One, I think that's really what we want to ... It's [inaudible], right? And two, I wanted to make up a new word.

TRIPTI SINHA: Terry, I agree. The work shouldn't be done here, but I think our effort here is to scope so that when it goes to the caucus, it's a well-scoped problem statement and it doesn't suffer from ratholing and scope creeping and so forth.

So, I think what you said is correct. Let's not do the work, but I think let's do the work in the effort to scope the problem statement.

WES HARDAKER: So, following on in my bullet one, if we phrase number two such that we say use the new metric to evaluate the Internet at large, they can't do number two until number one is complete. That provides some scope.

TRIPTI SINHA: Your number one is actually I think Liman's number 2a, what does underserved mean?

LARS LIMAN: Did we come to a conclusion or how to replace geo or I'll just take a stab at it and see what happens? Is anyone opposed to removing the phrase geo from this? Okay. I'll try to come up with something. Any other comments?

UNIDENTIFIED MALE: [off mic].

LARS LIMAN: I at least have some and I will come nag you all. Thank you.

TRIPTI SINHA: Okay. Thank you, Liman. Are you next?

WES HARDAKER: Yeah. Actually, next on the agenda is studying resolver behavior, but if I may interject one more document before it, because it should be quick. We met yesterday about the KSK, responding to the request by the board to talk about the KSK rollover plan. We did manage to keep our scope down to just considering the root server system's scope and not expand into other directions. The document and corresponding why we're writing it document are nearly done.

So, my question to this group and our question, as we discussed it yesterday, is in the interest of time, rather than send it to RSSAC for a week review before sending it to the caucus, we'd like to send it to both the full caucus and RSSAC at the same time to shorten the timeline that the caucus has to work on it. Or, excuse me, extend the times that it needs to be approved by end of July-ish in order to get submitted by August 6th. So, we're likely trying to do an e-vote at the end of July would be my guess. But, does anybody object to sending it to the caucus and the RSSAC body for the first read at the same time? Brad?

BRAD VERD:

I have no objection to that question, but we do have a work session on that tomorrow. Correct? We have 90 minutes allotted to working on it tomorrow. So, we would complete that before it got sent anywhere, correct?

UNIDENTIFIED MALE:

Yeah. It was a working session scheduled 5:00 PM, but it was mostly for the work party.

BRAD VERD:

Right. So, do we have a session in the full agenda?

WES HARDAKER: On the agenda tomorrow, 10:30 to noon RSSAC work session number four, RSSAC 000 we would cover and KSK rollover plan response. That's tomorrow.

BRAD VERD: Okay. So, that means for this work party, we had better wrap it up tonight so that we can send it to people, so we can review it by tomorrow. Okay.

LARS LIMAN: But, the original question for Wes, does anyone object to sending it to this group and the caucus at the same time? That's the question.

WES HARDAKER: And if we can get most of it done ... We can get the living document done by tonight and send a viewable link to the caucus and RSSAC [inaudible] and mention that it will be discussed in the meeting and that we encourage caucus participants to call in.

TRIPTI SINHA: Does anyone object to that? Hearing none ...

LARS LIMAN: That might be too short.

BRAD VERD: Again, I think that's a really short time. To me, why wouldn't we just have the discussion tomorrow here? It's an open session. People could be here and listen just like it's normal, and then we can send it to the RSSAC and caucus then. It's just a suggestion. I'm not opposed to sending it to both parties. I'm just talking about timing. That's all. Because you're going to send something tonight that's going to get changed tomorrow and resent again. And by the time somebody might see it for the first time, there will be potentially a couple different iterations of it. That's all.

WES HARDAKER: I think we're mostly in agreement. Before the discussion at 10:30 tomorrow or whatever time it is, I'd like to get a copy to everybody so that they can read it. And because this is an open session, I'd actually like to suggest that even the caucus members that want to participate on this timeline can engage.

FRED BAKER: Silly question. If you send it to the caucus, we all get it, don't we?

WES HARDAKER: Correct.

FRED BAKER: So, what's the difference between sending it to the caucus and sending it to the RSSAC and the caucus?

WES HARDAKER: None. Good point.

BRAD VERD: When you craft the e-mail for whatever is going to go out, make sure you give a context that, given the timeline and the response time, we think it prudent that this get out there as soon as possible so we get feedback. You know what to say.

WES HARDAKER: Thank you for reminding me. I'll make sure I do that. One second before I go on.

BRAD VERD: This will be a deviation and normal procedure. I don't want somebody in the caucus who is not used to getting e-mails about this stuff to be like, "Oh, I want to sign up for that work party." No, no, no. There's no work party here. We're moving full steam ahead.

WES HARDAKER:

Alright. So, that brings us on to the studying resolver behavior new work party. As I think we talked about at some point in the past with Andrew's survey and the work that we were going to put forward, one of the other big ones, there was actually two different items in that survey that talked about resolver behavior and we decided to merge them and put them into one study about – if I can stare at the right screen. Studying modern resolver behaviors.

Resolvers continue to change over time, and there's a lot of unknown with respect to how recursive resolvers pick authoritative servers, how they switch between them, how they decide which root servers to use. Do they treat the root servers differently? There's a lot of analysis that can be done there.

So, the link that I shared earlier, and it is also in the RSSAC comments, the Adobe Connect comments, is the conclusion of what a number of us put together – and I'm forgetting everybody that worked on it.

In my opinion, it is ready to go in terms of being able to send it to the caucus. So, we're asking for final comments from RSSAC about it. I don't know if anybody has time to read it. The proposed work party leader is myself right now, and I think Fred agreed to be shepherd. Is that what we finally ended up at? Okay.

Then, I think if this statement of work looks good, I think we need to vote on it in July to make it actually a work party.

TRIPTI SINHA: Can I have the link to the document? I'm sorry.

WES HARDAKER: No, please, that's why we're here. Yeah. So, do we want to take five or ten minutes? Yeah. It's the first link in the Adobe Connect chat messages. I also sent it in e-mail sometime this morning, I think. So, we'll take five or ten minutes for anybody listening, so people can read it. It's short.

Feel free, in my opinion ... So, I just changed it from editing to suggestion. If you want to make changes to it, wording or whatever else, feel free to do it live. Go ahead.

So, I think hopefully everybody has had enough time to at least skim through it. I am not going to go through. We won't iterate through the wording. I don't think we need to do that, so instead, I'll suggest ... I'll go read wording suggestions later and approve or whatever. In the meantime, does anybody have comments that are worth of discussion or things they want to bring up? Rather than me looking through them, you tell me if you have ones. Brad?

BRAD VERD: I have just a couple of clarifications. Number one, analyze DNS resolver network traffic and behavior. This has been done numerous times, right? This is done all over the place on how the different resolver softwares kind of choose which server to talk to.

WES HARDAKER: That's a good question. I think part of this will be collecting those previous analysis and seeing if they're up to date. Resolver software has actually changed a lot, especially in the last decade. In my opinion, one of the most useful starting things to do will be to go collect all those, summarize them, and then sort of create an aggregate report, as well as identifying missing holes.

I'm not sure ... I think that they were all done in different ways. I think some have looked at source codes, some have actually tried to study traffic. I'm not sure that we have a good consensus on ... I think we have a good consensus on unbound and BIND. I don't think we have a good consensus based on traffic that we see in other places. So, I think an in-depth survey of DIDL, our recent DIDL would be good. But that's a good question.

BRAD VERD: While I think that's interesting on how the software does it, I think what you'll find when you look through the different ways its been done and the different times its been done is that it changes every implementation – not every implementation, but it changes. Let's just leave it at that.

I think maybe the question that we could be asking is: should it change? Maybe that becomes one of the recommendations or a possible outcome type of thing.

WES HARDAKER: That is actually listed in there. I know because I edited it this morning or last night. Number six. If it becomes apparent that further work on the DNS protocol is needed to better optimize implementations and deployment. I [inaudible] properly. I was trying to make that bullet say both. We might want to change implementations. It has nothing to do with standards. And we might need to change standards and maybe splitting it into two bullets [inaudible].

BRAD VERD: The other piece is, again, putting context around. My guess is that whatever you come up with, number one, it's only good for this time because it will change.

WES HARDAKER: So, let me ask you a return question. I guess two things. Are you arguing that we shouldn't be doing it, that we don't have a need, or you just want to make sure it's properly scoped to say it's going to be time-limited and useful as well?

BRAD VERD: I'll return the question with a question. As a researcher, would you like it to be predictable or not? It seems to me like you'd want it predictable, meaning not to change. You'd like resolvers to ... Maybe each resolver does it differently, but at least they keep it that different in perpetuity so it's predictable. Or not. I don't know. I feel like number one has been done numerous times throughout the years. Back when there was a time when BIND went to A first, right? Then we finally got that fixed.

WES HARDAKER: I think my point is that it's time to do another ... To be fair, in full disclosure, I currently have a research grant proposal out to do this type of stuff, to study it. I have a colleague at ISI that's planning on joining the caucus work party because she and I are the PIs for it. Regardless of whether we get that research funding, we would contribute anyway. If we get it, we'll be better off.

In our view, the modern – and that’s why modern is in the title. It seems like it’s time to do it again. You’re right that it has been done in the past.

BRAD VERD:

Me personally, personal opinion, I think the one through four are cool and it comes up with the data you need, but I think, for me, what’s interesting is six and potentially seven. Do we have a recommendation that would create a standard so that this behavior either stops, and if so, why, type of thing? I think that, to me, is what ...

WES HARDAKER:

No, that’s a very valid point and I agree completely that the biggest outcome will be what can we change? Just the analysis, it’s like who cares if it doesn’t lead to improvement.

BRAD VERD:

And that’s kind of what I’m saying. The analysis has happened numerous times and nothing has really happened other than trying to understand what was going on at that moment in time. I don’t think we’re trying to figure out what’s going on at this moment in time other than the behavior. What do want the behavior to be?

WES HARDAKER: So, would you suggest changes to the document in structure or form or content in order to bring about that distinction in a way that you would prefer?

BRAD VERD: I'll think some more.

TERRY MANDERSON: Maybe I can help here. Yes, you're right that the last two points are the key things here, that the first parts add rigor. So, maybe we focus on the last two points as the headline, and then to get to those last two points, we expect the work party to do this in terms of to focus on these things.

BRAD VERD: Here's my fear. We do one through four. Five, six, and seven are hard. So, we're not going to ... That's going to take a lot of work.

WES HARDAKER: So, let me propose a restructuring. How about we list ... I'm not even going to look at the numbers. We list the conclusions we want, which is what changes to the systems and implementations do we need? Then we put the ... In order to do

this, you will likely need to do these and put the analysis afterwards. Okay, I will restructure that later. Matt?

MATT LARSON: I have a couple of comments. To my knowledge, there is no IETF document that specifies server selection. I'm pretty sure there is not one and that it's literally language in 1034 – no, it would be 1035 – [inaudible] choose the best server. So, we're only about 25 years overdue for such a document.

UNIDENTIFIED MALE: I agree.

MATT LARSON: So, I thought about it over the years, but never done anything about it. And I think whoever writes that can be prepared for a long slog ahead of them. I disagree that five, six, and seven are harder than one through four.

BRAD VERD: Just to be clear, I'm not saying they're harder. I'm saying we've had a couple of different work parties now that have come up and said – not really answered the questions that were asked.

MATT LARSON: Fair enough.

BRAD VERD: What I don't want to have happen is to go through this and have six and seven not answered. That's what I was trying to say.

MATT LARSON: My other comment is I'll second the comment Duane made first, which is I think this needs to be more than analyzing code bases. We did similar work—

WES HARDAKER: I'm sorry, more than analyzing what?

MATT LARSON: Code bases. That you have to actually simulate this in a lab environment. We did this similar work on server selection in Verisign labs, I don't know, seven years ago maybe, and we found things that I don't think you ever would have found just by code inspection. You had to simulate them. I also suggest that be made more clear.

TERRY MANDERSON: I'll echo that as well. One of my guys did a small lab experiment and we were floored after looking at some of the code. It just

didn't behave the way the code expected. There may well just be bugs, but it's interesting.

WES HARDAKER: I think if you spend any amount of time talking about [Mark Andrews] about corner cases, it's like why are you re-querying this when it should have been in your caption? It's like, well, it had to be [inaudible] this way because of this, this, and this. I think you're 100% right. So, feel free to, again, modify text accordingly.

UNIDENTIFIED MALE: And at the risk of scope creep, what would be great is if the work parties are going to go to all the [inaudible] simulate it to do it in a repeatable way, so that there's a way that when there's the next version of [inaudible] not whatever, that we've already got the test infrastructure so that somebody can build a few VMs and test it.

WES HARDAKER: Yeah. Reproducibility would be a good thing. Thank you.

BRAD VERD: That would feed the tools, also.

WES HARDAKER: Yes. It would feed right into the tools. I will create an account. The one thing that I think will have a struggle with is where to define certain outputs. The IETF has traditionally shied away from is this an interoperability problem. So, something like server selection or how do you rank servers? That's entirely possible that that has been left out of the past 25 years because it's not needed for interoperability, even though it might be needed ... It might come out as best practice or something. We might have to have scope it appropriately.

UNIDENTIFIED MALE: Yeah. I'm thinking [inaudible]. Yeah.

WES HARDAKER: Alright. I think that closes those two comments. Anybody else have things they want to bring up about the document?

UNIDENTIFIED MALE: Is the intent still to vote on the 10th? So I can add that to the agenda.

WES HARDAKER: I think so. What's the timing? I need to look at a calendar. I have to have this done by a week before the ... Is it the 10th? Okay. So, I will wrap this up and send it out as a, hopefully, stable

document. I'll give it one last round. I'll send it to RSSAC to say, okay, last call on changes. Then we'll have it stable by the 3rd so that people can vote on the 10th.

TRIPTI SINHA: Done, Wes?

WES HARDAKER: Yes.

TRIPTI SINHA: Alright. Very good. We've got two documents to be delivered on July the 3rd as stable documents. One that you're leading and the one that Liman is leading. Correct, Liman? Right? Okay. So, we will plan to vote on, at a minimum, two documents on the 10th.

Anything else on the agenda for this morning, late morning? Then I believe we're adjourned if we have nothing else. Do we have anything? Anyone? Any last comments? Alright. We are adjourned. Thank you.

[END OF TRANSCRIPTION]