ICANN67 | Virtual Community Forum – At-Large Policy Session - DoH/DOT - Threats and Challenges
Tuesday, March 10, 2020 – 13:00 to 14:30 CUN

PAUL HOFFMAN: Hi. This is Paul Hoffman. Just doing a quick soundcheck before I give a presentation later in this session.

YESIM NAZLAR: Hi, Paul. Welcome [inaudible]. Thanks so much for the audio check. Loud and clear. Thanks so much. And welcome to the section.

JOANNA KULESZA: Hello. This is Joanna on the audio bridge. [inaudible]

YESIM NAZLAR: Hi, Joanna. Welcome. Loud and clear. Thank you.

JOANNA KULESZA: Brilliant. Thank you so much.

UNIDENTIFIED MALE: I think everyone is just silent.

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.

[JULE HEDLUND]: Hello, everyone. Welcome to the call. We are going to collect Holly Raiche. We will begin in a few moments. Thank you all for your patience.

YESIM NAZLAR: Hi, Holly. Welcome. This is Yesim speaking. Would you like to do a quick audio check so we can start the call?

Holly? [inaudible]. Are you able to hear me? Can we please do a quick audio check, as you'll be moderating this call?

HOLLY RAICHE: I'm fine. I'm looking at the wrong screen. Is the screen up?

YESIM NAZLAR: Holly, currently I'm showing the expected standards of behavior. If you're ready, we're going to start the recording. I'll read the reminders and then [inaudible]—

HOLLY RAICHE: Who's talking?

YESIM NAZLAR: Okay?

HOLLY RAICHE: Yesim, who's talking now?

YESIM NAZLAR: Apologies. Looks like we're having some technical issues. If you could please bear with us for another minute or so, we'll get this sorted as soon as possible. Thank you.

Hi, Maureen. I see you're already connected to Zoom. Would you please check your audio so we can get started and maybe I can hand the floor to you.

MAUREEN HILYARD: Hi there. I'm just getting myself organized here with a [inaudible]. Sorry.

Okay. Thank you. We can get started.

YESIM NAZLAR: Sure. Good morning, good afternoon, and good evening to everyone. This is Yesim Nazlar from At-Large staff. Welcome to the ICANN67 virtual meeting and the At-Large policy session—DNS-over-TLS and DNS-over-HTTPs: Threats and Challenges—(virtual session) on Tuesday, the 10th of March, 2020, at 18:00 UTC.

The Zoom room audio is in English. In order to access the French or Spanish audio, please join the French or Spanish training via the link on the main ICANN67 website.

All details were sent out on the ALAC announcement with all relevant links. Details for these connections can also be found on the ICANN67 At-Large wiki agenda pages.

We will not be doing the roll call today for the sake of time. ALAC members, [inaudible] leadership and liaison's attendance will be noted.

If you would like to ask questions or make a comment in English, French, or Spanish, please type it in the chat by starting and ending your sentence with Question or Comment, and please keep them short if possible. French or Spanish questions will be translated into English and read aloud by our remote participation managers: my colleague, Claudia Ruiz, and myself, Yesim Nazlar. Staff will put periodic reminders of this process in the Zoom room chat. So, if you are in the Zoom room and wish to speak, you may also raise your hand. Staff will manage the queue.

A kind reminder to please state your name when you speak not only for transcription purposes but also for the interpreters to identify you on the audio streaming. Please also speak clearly at a reasonable speed to allow for accurate interpretation.

Finally, this session, like all other ICANN activities, is governed by the ICANN expected standards of behavior. I'll put to the link in the chat to those standards for your reference.

Without further ado, I would like to now turn the floor over to Maureen Hilyard, ALAC Chair. Over to you, Maureen. Thanks so much.

MAUREEN HILYARD: Thank you. Thank you, everyone. I hope you can hear me well. I'm just standing in at the moment until Holly organizes her phone connection here. The session, of course, is related to the DoH/DoT threats and challenges, a topic that Holly Raiche has organized for this session. I hope that they will let me know as soon as she becomes available so she can actually introduce it. It's not an area of my expertise.

I do note that we have our guest. We have Paul and, I assume Rod, available as well. I just want to welcome—and I note Barry is listed as well. I hope that everyone is probably a little bit more prepared than I am for stepping in at this particular point in time.

If you can just bear with me so I can just check that Holly is available now. Thank you.

I'm not getting any messages here. Okay. To start off with, perhaps we could have an introduction. I know Holly has had some discussions with Rod [inaudible]. Could we have an introduction, perhaps, Rod? This is an issue that has actually been discussed within At-Large before, anyway, but perhaps a summary of what discussions have been taking place in relation to this topic. Would that be possible?

ROD RASMUSSEN: Hi, Maureen. We could certainly have a conversation around that or some discussion around that as we're waiting for Holly.

I believe we had a … I've seen the schedule. I'm rambling because I'm remote. I'm on my computer. I believe we had a schedule where I think Paul was going to run through a presentation, and then I believe Barry was going to run through a presentation as well.

Paul, is that your recollection? Am I correct in that?

PAUL HOFFMAN: Hi. Actually, a little bit before this meeting, it was decided that I would be giving a presentation and that there wouldn't be a formal SSAC presentation but you could informally talk about what's coming up in the future and such.

ROD RASMUSSEN: Okay. Because we had prepared some slides. I'm—

HOLLY RAICHE: Folks, can I actually talk now? Can people hear me?

ROD RASMUSSEN: Yes.

MAUREEN HILYARD: We can hear you.

HOLLY RAICHE: I'm listening to Paul's presentation, and it's fantastic.

PAUL HOFFMAN: Holly, I don’t believe we were listening to my presentation because I’ not giving it. We had an issue earlier where it sounded like you were listening to a tape of a previous presentation.

HOLLY RAICHE: Oh, okay. Yes, I was. I’m happy to start. I really apologize. I had it down for 5:30, so scream at me. Can we start?

ROD RASMUSSEN: We’re not going to scream at you. Go ahead.

HOLLY RAICHE: I’m ready to start. Okay?

ROD RASMUSSEN: Go for it.

HOLLY RAICHE: Okay. What I’m listening to in my headset is Paul talking. So can I hear this session or do you want me to talk?

ROD RASMUSSEN: Holly, did you have an introductions that I wanted to do?

HOLLY RAICHE: Yes, I do.

ROD RASMUSSEN: All right. So why don't you go ahead with those? I'm not sure where you're getting Paul's voice from. He's not talking on this call right now.

HOLLY RAICHE: Okay. Do you want me to start now?

ROD RASMUSSEN: Yes.

HOLLY RAICHE: Okay. Folks, this session is going to be looking at DNS-over-TLS, which is DoT, or DNS-over-HTTPS, which is DoH. That means DNS over the transport layer or HTTPS, which are both secure systems. That's the introduction that's in the agenda, but it doesn't really explain what either mean.

Looking at some of the previous sessions, a lot of the focus of ICANN has been on privacy and what personal information should be made public, should be made accessible, how, and to whom. But nowhere near enough attention has been paid to an equally important privacy issue for ICANN. In this case, the information is about people—about where they go onto the Internet, about who they contact on the Internet, and about what they do on the Internet. For law enforcement agencies, this information is called meta data. It's as important, if not more so, than someone's name and contact details, yet this issue has really gone under the radar, so to speak.

So Paul Hoffman is here to explain all. Just a word of background. Paul Hoffman joined the ICANN staff in 2015 and currently serves as the principal technologist. Serving as the senior technologist, Paul is responsible for improving ICANN's technical capabilities and stature both internally within ICANN as well as externally within the Internet community.

Over to you, Paul.

PAUL HOFFMAN: Thank you, Holly. I'm going to give a mediumly-brief conversation covering some of the aspects Holly just talked about because they are quite important. Then I believe we will have some discussion from the folks at SSAC—Rod was already speaking—and there'll be plenty of time for Q&A after that.

Let's just go ahead and go to the next slide, please. I'm the principal editor of a document that we just recently published. In fact, since not everyone can click on their screens, I just pasted into the chat a direct link to the PDF of that document. The title of the document and this discussion is "Local and Internet Policy Implications of Encrypted DNS." As Holly as just saying, there are a lot of policy implications of people seeing each other's meta data and things like that, knowing what you are looking for. To be clear, these are both positive and negative. So we'll be talking about positives and negatives in this.

SSAC will be coming out with a very different document in the near future. Rod will talk a little bit about that. It goes in different places.

For this document, basically there were four broad topics. One is filtering and monitoring in the DNS and why people do it and the implications of that, then the policy implications in general of being able to encrypt DNS traffic, who are the interested parties, and then some of ICANN's positions.

Let's just keep moving forwards. I think it would be best to hold questions until the end, even if they're technical questions. It turns out that the technical side of this discussion is much less interesting than the policy side.

Next slide. This is the one graphics slide that is obligatory in every DNS presentation. Who are the participants in the DNS? In this case, we are talking about the technical participants, not the people. The technical participants are always working on behalf of the people. On the left hand, you have the DNS stub client, which is the thing that's running on your phone or in your computer or, in some cases, in your browser. Slightly to the right of that, you have recursive servers, which are systems that work on the behalf of the users. Then you have the obligatory giant cloud of the Internet. On the righthand side, you have the authority servers, which are the ones who really know the answers to the DNS questions that you're asking.

Many of you have seen slides like this before. The most important part of this slide is that all of those gray arrows currently are unencrypted. The question of encrypting DNS is all about, at this point, encrypting just those two arrows on the left—the two ones going from the stub client to the recursive. There are opportunities to do more encryption later, but, for our discussions today, it's just the question of encrypting from the stub client to the recursive server.

Next slide. Let's talk a little bit about where people are talking about doing encryption. As I just said, right now, we're biting off the first piece, and that's basically because most of the interesting policy implications are in fact about encrypting from the end user to that first [hop]. Beyond that, the need for encryption, the value of encryption, goes down. So right now we're only talking about the left.

Now, stub resolvers actually appear in a couple of places. Up until very recently, everyone assumed a stub resolver was part of your operating system. Recently, applications (particularly browsers) have been implementing their own stub resolvers. So, on you sitting on your phone now, everyone would have assumed you had one stub resolver that was part of the operating system of your phone. Now you might have a few of them—some in the browsers, some even in games or any other applications.

Let's skip ahead here because really what we get into now is the how, not the where. So next slide, please. Sometimes the IETF does it the way everyone expects and comes up with a standard for the way things should go and how we're doing things. Sometimes the IETF does it slightly wore, which is to have multiple ways of doing the same thing. I say that jokingly here because I'm actually co-author on both of these documents. So I take partial blame.

There are two standardized protocols for encrypting DNS currently. One of the is called DNS-over-TLS. Almost everyone refers to that as

DoT. The other one is called DNS-over-HTTPS, which is called DoH or, for those of you who are fans of the cartoon The Simpsons, they might call it, "D'oh!" There are links here for how you can read those. Basically, the most important thing for a conversation like this, which is really about policy and implications, not about the technology, is that they really, really are similar but they do have some difference that are very important to network operators. We will focus only on the differences before we get into the policy section.

Next slide, please. For DNS-over-TLS, which is the first one which came out, the stub resolver sets up a TLS session with the recursive resolver that is virtually identical to the way that your browser starts up a TLS session when you go to a webpage that also is encrypted. So TLS is the security protocol. It works in lots of contexts, and it works fairly well in this context.

Now, one of the things that many of you have seen over time is that, in TLS, authentication is very important. That is, if you attempt to go to a website, and that website can't be authenticated in TLS, the browser is going to tell you that and try to convince you not to go. In DNS-over-TLS, the authentication of the recursive resolver is somewhat optional in that we're just trying to get whatever privacy we can, but it is needed if you are wanting greater assurance. So DNS-over-TLS is actually fairly easy to set up.

Next slide, please. DNS-over-HTTPS, which is the one that came a few years after DNS-over-TLS is where a stub resolver starts an HTTPS session, meaning it is literally going to be sending a request looking just like it would have if it were a web browser. So this is using full HTTPS, and it uses all the semantics that the web people have done for us.

You can look at the history of the Internet with respect to this in two ways. One is, “Well, the DNS came first, and therefore it’s more primary.” Or you can look at as, “You know, the web people did orders of magnitude more work to get this going than the DNS people did. Therefore, it’s more important.” That’s really the genesis of why we have two different ways of doing DNS encryption. For example, if you’re doing DNS-over-HTTP and you’re using Version 2, it actually allows things that DNS-over-TLS (or DoT) doesn’t allow.

So that’s really the big takeaway here: yes, there are two, but the new one has some more examples. Again, technical advantages are not really what we’re talking about today because we really care much more about the policy stuff.

Next slide, please. The document that I was the primary editor on has these seven thing as the major policy implications of doing encrypted. We’re going to talk about these in the coming slides, but, just as an overview, when you do encrypted DNS, you get increased privacy of the user’s DNS traffic and you get increased insurance of the user’s DNS traffic. If you just looked at this slide and looked only at those two bullets, you’d say, “Oh, good. Encrypted DNS? That’s a very good thing.”

For the next three bullets, you’ll notice that we use the word “circumvention.” That has a fairly negative connotation for a very good reason. Encrypted DNS is often used to circumvent the DNS filtering policy that someone already has or their local policy or even the policy that is mandated by the governments. So, if you looked at this slide and you skipped the first two bullets and only looked at the middle ones, you would say, “Oh, my. Encrypted DNS is pretty bad.”

The last two bullets is also important. One is that you have the possibility of getting unwanted centralized of DNS resolution. There is wide disagreement about whether that’s bad or good. Also it’s speed of DNS responses. Again, that’s a technical thing. We’re talking really milliseconds there, so it’s not so important for our discussions today.

Next slide, please. Let’s go through what I labeled as the two positive bullets. Privacy is generally good. We are trying to give end users more privacy. If many end users are using the same computer, and that computer is using encrypted DNS, all of them get that kind of privacy. Basically, this is privacy from people who are snooping the Internet—not looking at their computer, not looking at their end thing—and going, “Oh, look. This person is asking this DNS query,” or, “This person is this active.” Things like that. So that has generally be considered good, and was one the main reasons why we started this work.

Also, using encrypted DNS also prevents an attacker from changes the responses. So, when you have an encrypted session, you are sure that that answers you’re getting are exactly the way that the host gave you. For example, an attacker cannot see if you’re asking for an address of how to get to a particular government website. They can’t jam in a wrong answer that would then fool you into going to the wrong place and possibly giving up private information. So DoH and DoT increase these things pretty similarly.

Next slide. So those middle three bullets were about circumvention. This is one of the reasons why so many people have become so concerned about encrypted DNS: basically, you have the issue of … I think I'm getting some echo from the background there. I hope you all aren't hearing it. Basically, this is about the issue of, if you are using a service that is purposely filtering or monitoring your DNS traffic … For example, one of the things that some ISP do is they look at your DNS queries and they see if you are making a query for something that would lead you to a malware site. They're going to stop that. They will either stop your query or they will send back an answer that is purposely a lie so that you don't go to the malware site.

This is actually more common in some regions, in some countries, than in others. Even within in region—for example, within the United States, which is where I'm from—some Internet service providers do this for free. Some charge a little bit of extra money. Many of them don't do it at all. But this is often a service that is given to increase the end user's security.

Also, some of this filtering is mandated by governments. There are laws in various countries that say that service providers have to do DNS filtering in order to prevent people from getting the correct address for certain types of sites. If you're encrypted DNS in a way that prevents that filtering, then the user is now no longer getting the benefit or the harm or whichever from those government laws.

Next slide. By the way, we are going at a fairly good clip here so that there'll be plenty of time to ask questions. I can vaguely see a lot of questions and answers going by over on the right. I'm glad that some of you are answering for each other.

There is also the question of unwanted centralization. Some clients that implement encrypted DNS can do it not only in a way that says, "Here we're going to encrypt what you would have been sending," but they actually also send it to a different place that supposedly will give you better privacy than what, for example, your ISP might have done. So, when they do that, they're doing it for what they consider to be a benefit, which is to give you even better privacy. But the net result is they might be sending all of those queries to a very small number of resolvers.

So there's a positive and minus here, the positive being, if you trust your software to be making these choices for you, it gets to do that. The minus is, if you don't know that it's doing it, it's going to be sending all your queries to a smaller number of places, and people might be able to use that in order to deanonymize you.

So, so far, this is only common in one place, which is in the United States, and it's only for people who are using one browser (Firefox), although that could change in the future. The different browsers look at what each other is doing all the time and such like that.

So the unwanted centralization discussion is very hairy. Lots of people are very concerned about it. So that certainly is one of the ones that we'll be discussing.

Next slide. You know what? Let's skip over speed of responses. That's a nerdy little thing. Again, we're talking a small number of milliseconds here.

Next slide. In the paper that I was editor of, ICANN actually took some positions. Since this is an ALAC call, many of you are much more aware of this than I am, even though I'm on ICANN staff: there's a difference between positions and policies. Positions are things that we can say that are not policies but are general things that are good in the world. One of ICANN's positions in here is, "Privacy is good." So increasing user privacy is considered to be a positive. It's not a universal positive, but it is definitely a positive.

Filtering in the DNS can also be beneficial. We have seen many places where filtering has in fact prevented harm from coming to people. Not all filtering is beneficial, so that's why this says, "It can be beneficial."

The third bullet gets a little bit trickier. The third bullet is that, currently, the applications and the operating systems—the things that are doing DNS queries on the benefit of the users—actually don't currently have enough information to actually make the kind of decisions such as, "Oh, this user is using an ISP whose filtering is good. Therefore, I won't touch [him]. Or, the flipside of that is "My God. This user is on a network where it's really pretty scuzzy. Therefore, I should step in and do something." Right now applications and operating systems just don't have very good information on that.

So the last position, which is really the basis of almost everything we do at ICANN, is that DNS data should be protected as well as we can for the benefits of users.

Next slide. This is the last slide I have. For those of you had seen the document earlier, we did an update of just a few weeks ago. So we're on Version 2. Basically, it's the same document, but there are some recent updates to this discussion, which also might come this session as we are talking now.

One is that Mozilla, the folks who do Firefox, is greatly expanding their program in the United States—and, again, so far nowhere else. In the first version of the document, we were talking about that Mozilla would be doing something in the future. Now they are starting to actively do it. They've done a whole lot of publicity on this. So it's reasonably well-known, but it is still considered somewhat of a lightning point.

Since we did the first document, Microsoft also announced that they are going to start doing encrypted DNS in Windows—that is, in the operating system, not in the browser—and that they—I'm hearing some background stuff there—will be doing it using DoH, not DoT, which has confused a number of people. That hasn't been released yet, but they're talking about doing it maybe sometime this year.

Really, the most important thing is—we'll see this on this call—in fact that network operators have really started to pay attention in the last six months, and they're getting much, much more involved in these discussions because users often think that they're getting their DNS from a network operator, either from an ISP or, if you work in a company that has your own IT department, they're you're network operator. Those folks have gotten much more involved and we're seeing them getting active in many places.

That was the last of my slides. Holly, I don't know how you want to continue from here. I'm happy to be answering questions. I bet Rod is as well.

HOLLY RAICHE: A few questions now. First of all, thanks very much. In the chat, we've had a request for a link to your paper. I think they're talking about … The last one I saw was, I think, February 24 as your Version 2?

PAUL HOFFMAN: Yes, that is. Here, I'll paste it in again now because it looks like the chat has been quite, quite active. So I'll paste the same thing in again.

HOLLY RAICHE: It has. So people would like to see the paper. It's only twelve pages. People can read it. It's very interesting.

Let's actually start. If people want to put their hands up. The first question really was from Thomas De Haan from the European Commission. He has had a couple of questions.

Thomas, are you on [the phone]? Can you unmute yourself and talk, or do you want me to start reading the questions off the chat?

Actually, I'll go back to the chat and start reading questions. Let me find him.

THOMAS DE HAAN: Hello? Can you hear me?

HOLLY RAICHE: Yes. Go ahead.

THOMAS DE HAAN: Okay. I had two questions. The first one is … Maybe the second one is more important. We heard from Chrome that they would deploy in March 2020, actually. They talk about the global [inaudible] deployment and they talk about [auto]—it's the same resolver—upgrade. So I wondering if somebody [inaudible]. Thank you. And [inaudible] DoH [inaudible] DoT, but I already got the answer in the chat. Thank you.

PAUL HOFFMAN: I'll do the second one first. If you got an answer, I'd love to hear it because I get such confused responses on the two of them, which, again, is somewhat surprising to me because I'm co-author on both of the standards.

Since this is ICANN, we care very much about names. We have a naming issue in your first question, which is Chrome. There are many things that Google has called Chrome. As far as I know, the Chrome browser in fact will not be deploying encrypted DNS at all. That is, they have made no announcements about them doing that.

They have made announcements, however, about them … In fact, they are doing encrypted DNS in the Android operating system. The way that it works in the Android operating system currently is that, as you connect to your resolver, they check if your resolver will also do encrypted DNS. If so, they will, without telling you, do an upgrade for you where they will just start encrypting the traffic.

So far, that's all we have heard from them about how they are handling this. They may in the future do things differently. We don't know, but, right now, the major thing that we know that they are deploying is that they'll do it.

Now, the Google folks have said that they are going to be doing this in Chrome at some point, but we haven't heard from them since then. We passed by when they will be. So I'm taking that as we don't really know. Both Google and Firefox have been very good about saying things at least a few weeks before anything is deployed.

HOLLY RAICHE: The next questions—actually, Jonathan, you had about two or three comments in the chat. Do you want to ask a question now? Because otherwise we're going to go to Dave as the next one. So, Jonathan, over to you.

JONATHAN ZUCK: Go ahead, Dave. I can wait.

HOLLY RAICHE: Okay. Dave, can you talk or shall I read out your question? I'm looking at the chat.

I'm not hearing anything. Paul, the question that Dave asked is, "What is the rationale for encrypting only partly the session between the DNS sub-client and the DNS recursive server? Why not the sessions between the DNS recursive sever and the DNS authoritative server? Is it not partial encryption?"

PAUL HOFFMAN: Let me preface that with saying that asking an individual about the rationale that other people might have is always a bad idea and it's also a very common thing to happen in any of these discussions.

The rationale for why I personally, for example, have not pushed so much for encrypting the whole channel is there is a lot less benefit to encrypting between a recursive resolver and the authoritative because a recursive resolver in general is responsible for the queries for dozens, hundreds, thousands, or possibly even millions of users. So the personably identifiable information gets a lot less. It's not zero, but it gets a lot less.

The second reason is, quite frankly, it is much harder for an authoritative server to figure out how much technical resources they have to put in to turn on encryption than it is for a recursive resolver. I believe we will end up in a world where all of the DNS could be encrypted, but, right now, the technical world has really been focusing on the first part because that's the part that most directly affects users.

Rod, do you want to speak to that? Because I think that that is actually part of one of the things that SSAC looked at.

ROD RASMUSSEN: I'm going to hand it over to Barry, who is one of the Co-Chairs of the work party, for questions on this. Just to give some background, we've been working on a paper on DoH/DoT for a while now. I think we're shared that with ALAC the last couple of meetings. We were hoping to have it published before the meeting. We literally just saw the copy, the PDF of the final version, within the last hour. So we will be putting it out there. We have some last-minute discussions about what was in the paper, and we had to do our SSAC magic to resolve all that, which did happen, which was great. But, unfortunately, we didn't get it out just before this. But it'll be available shortly.

We also had a presentation. I'm not sure what the plan is. If we don't have time for that, that's fine. I'll turn it over to Barry Leiba to drive the SSAC response here, if that's all right. Barry?

BARRY LEIBA: Sure. Hi. I don't have much to add beyond what Paul said. The SSAC report does not cover the idea of encryption between the recursive and the authoritative because the protocols currently don't specify that. But we did discuss it in the work party.

I agree with Paul's answer. It's basically that the concentration of private information is less between the recursive and the authoritative. So it was easier to get this deployed— DoT, in this case— without going on that end of it.

Nothing else to add.

If I may, I saw a couple of other questions roll through that I'd like to address. One of them in particular that we do talk about a little bit in the paper from SSAC is DNSSEC versus DoT and DoH. And there've been a lot of answers rolling through on the chat.

The bottom line is that DNSSEC addresses one aspect of securing DNS responses, and the encrypted protocols look at a different one. The encrypted protocols are preventing snooping on the pipe. They allow servers to communicate privately. But a server that has bad information can still transmit that bad information even though it can't be looked at or modified [en route.] DNSSEC is designed to prevent the bad information. It authenticates the response themselves—the response data—with signatures. So they're complementary, and both protocols are still very important.

The other one that I wanted to answer that I saw rolling across was, why do browsers want to do DoH rather than Dot? The answer to that is that browsers and other web-based applications already have the infrastructure built into their programs to deal with HTTP and HTTPS connections and all of the networking aspects that that involves. It's easier for them to implement DNS over that than having to use a different port and a different protocol to talk and have to put different networking software into their applications.

I hope that answers those questions.

HOLLY RAICHE: Thanks, Paul. It think I probably would have added, if you go back to the very beginning slide you have, which is, "Where does this happen? It happens in the lookup," DNSSEC is really about checking that you're going where you think you're going, and [inaudible] coming from where [he] thinks it is.

This technology is about protecting the lookup. Different thing, different security, different problem solved. Am I wrong on that?

PAUL HOFFMAN: No, you are not wrong, but what you just said is only part of the answer. I would love to hear also responses from SSAC on this because SSAC is very concerned about authenticity of the DNS data.

One of the issues that we've always had is that—you can check the veracity of a response you get with DNSSEC (anyone can check the veracity) so, when DNSSEC was originally developed, everyone thought, "Oh, okay. Well, the user will check that"—today, basically no users check that. They just trust the resolver that they're talking to to check on their behalf. This lead to two problems. One is that you as a user don't actually know whether the resolver that you're talking to is bothering to check for you. The other is that you don't have a way easily of checking yourself without jumping through some technical hoops.

So the whole idea of trust in the DNS—trusting the answers that you're getting—is a lot less clear than it is in the web. In the web, you look up. You see a little lock or not. We can argue about that forever. But there is more visible feedback, and it is easier for you to say, "Why am I trusting this?" than you can in the DNS.

So, yes, you are correct that DNSSEC answers some of that, but only if you actually do the DNSSEC. As we've discovered, almost no end users do it. Actually, fewer than 30% of queries that go to resolvers do it. And, even when they're doing it, very few places use DNSSEC to sign their answers. So you can go to Google, and Google is not signing. So does that help you or not? It's really not clear.

HOLLY RAICHE: Excellent. We've got a couple more questions, and then we're going to go over to Barry. The first one I'm going to mispronounce, and I apologize. It's Gangesh Varma. "What has been the responses from law enforcement agencies across different jurisdictions to Dot and DoH?”

PAUL HOFFMAN: That’s not a question for me.

HOLLY RAICHE: Is that a question for Barry?

PAUL HOFFMAN: Or possibly for other people on the call who represent law enforcement.

BARRY LEIBA: Certainly not a question for me.

HOLLY RAICHE: I can understand. Barry, do you want to take that one?

BARRY LEIBA: No. I just said no; that this is not a question for me, either.

HOLLY RAICHE: Okay. Well, we’ll leave that one hanging. The next one is from Joanna. “How does the encrypted DNS requirements in local policy fill into the discussions on Internet fragmentation?”

PAUL HOFFMAN: I can take that. The question here of fragmenting the Internet, and by having local policy, is one that is very central to this. The short way that many people describe this is, "My network, my rules."

Now, how I connect my network to the Internet are on shared rules, but, once something is on "my" network, whether mine is my house, the organization for whom I'm IT, or all of my customers, even though I am an ISP, that gets a little bit tricky. Throughout the history of the Internet, we have silently if not explicitly allowed you to make rules. Things like encrypted DNS that can circumvent those rules might lead you to having an unexpected local policy. That is, you had a local policy that was working, and now it's not.

HOLLY RAICHE: Okay. We've got hands up. We've got Jonathan. Then we've got Gabriel. I was also going to read out one of his questions. Jonathan, go ahead.

JONATHAN ZUCK: Thanks, Holly. There's been a bunch of discussion in the chat about the dangers associated with centralization of the data. This is in large measure because of the small number of browsers. It's an oligopoly, if you will, of browsers that could lead to the centralization.

I guess my question is, they have the capacity to choose a single resolver now. That power isn't enabled by encrypted DNS. It's just something that might be prompted by it because of the unavailability of alternate resolvers, right? Am I missing something? It seems like, as soon as more server resolvers were capable of handling encryption, we'd end up back with the same kind of distribution curve [of] which resolvers were used.

PAUL HOFFMAN: I'll try to answer that, although, again, I speak only for as few people as possible, and you're asking a question that basically is a question that is aimed at folks at Mozilla. That's fine for now, but we might get more people doing it. Let me just take—

JONATHAN ZUCK: But the technical part of the question, Paul, is that they had the power to choose a resolver for you absent any implementation. That was the objective question, right?

PAUL HOFFMAN: Yes. Exactly right.

JONATHAN ZUCK: The power to centralize this data has existed all along.

PAUL HOFFMAN: Correct. That's exactly where I was going. Thank you. We are seeing much more discussion on this today, but the ability for a browser or any application on any of your devices to be doing what people are objecting to has been around literally for 20 years. The fact that

Mozilla told people that they were doing this is what got a lot of people's attention. But it is not new.

The other thing that's important here—again, I'm going to try not to speak for Mozilla—is, if you folks read the document that I wrote, I am actually speaking for them in that document, but I'm doing it a lot more carefully than I'm doing just off the cuff here. They are very clear to say that the recursive resolvers that they will be sending your queries to they have vetted and that they have strong restrictions on them. That's why the number is so small.

So, Jonathan, going to your question of, "Well, shouldn't this number grow?" the number grows only as rapidly as Mozilla wants it to grow. I'm sure that they would love to have a larger list because people would give them less grief about it, but they have standards for who they are willing to send your queries to. So far, they only have two. In fact, in Version 1 of this document, they only had one, who they would trust to live through those privacy restrictions.

It's very similar to the way that your IT department works. I work for ICANN, as people know, and ICANN has an IT department. ICANN's IT department has decided that they are going to run their own recursive resolver for the users at ICANN. Many of you on this call work in companies where the IT department has said, "Oh, my God. That's too hard. I'm going to outsource this," meaning I'm just going to send all the queries to some other open recursive resolver, often one run by Google, sometimes run by other people. There's lots of open recursive resolvers that would do this. That's a decision that somebody has made about your privacy. And it's very likely you don't know what that decision is.

HOLLY RAICHE: Excellent answer. We've got a few people in the chat. Gabriel, you've put a question into the chat. Do you want to ask that question now?

GABRIEL: Yes, please. I'll just read it out. I have a very slight addendum to it. The question is, does the unwanted centralization that was discussed lead to an increased risk of perhaps a large DDoS attack or denial-of-service attack targeting those servers, possibly leading to a largescale, worldwide web outage, much like what we saw in October of 2016 with the Dyn DNS attack.

PAUL HOFFMAN: Barry, do you want to take this? I suspect that this is something that you folks want to cover more in SSAC.

BARRY LEIBA: I'll talk about it. I think the bottom line on that part of it—the DDoS stuff—is that the people running the largescale recursive resolvers are aware of the issues with DDoS attacks have a lot of infrastructure in place to thwart or mitigate DDoS attacks. It's part of their business—the one run by Cloudflare, for instance. Cloudflare is a content delivery system, and they're well-prepared for this. I think we've come a long way in the couple of years since we've had the Dyn issue. A lot of people have learned a lot about thwarting these attacks.

So, yes, the more centralization you have, the fewer vulnerability points there are to attack but also more robust those points can get. So it's a balance.

GABRIEL: Thank you. That addresses it. If I could ask one small follow-up, is there anything inherent to the protocol itself that will switch to an unencrypted channel if the encrypted channel goes silent due to attack or otherwise?

BARRY LEIBA: Not in the protocol. The application itself would have to see that the resolver it was using wasn't working and switch over.

GABRIEL: Thank you kindly.

HOLLY RAICHE: Thank you. Alan, you're next.

Alan Greenberg?

ALAN GREENBERG: Ah, thank you. Someone just unmuted me. In this session and similar ones, we've heard a lot about words like "trust" and "privacy." What doesn't come clear really is it's not really a matter that one of these protocols or one of the methodologies is more trustworthy or more private than the others. It has a lot to do about who you trust. Yes, the traditional open DNS is snoopable by someone looking at the network. On the other hand, DoH or DoT is snoopable by someone running the resolver. Do you trust them? There was just some comments that Mozilla has only vetted certain providers because they trust them, but it really is a matter that all of these are vulnerable if you don't trust wherever the data is available. So you can't be traced by someone sniffing the network. On the other hand, the resolver operator or whoever is encrypting the questions for the resolver has the ability of tracing everything. If you don't trust them, then you're just as vulnerable.

So it really is a matter of, who do you trust? Not, is it better or worse? Thank you.

BARRY LEIBA: I'd like to take that one, if I may.

HOLLY RAICHE: Okay.

BARRY LEIBA: Because we do cover that in a good way in the SSAC paper. The first thing I'll say before that is there are a whole bunch of issues and solutions involved in making the DNS or DNS activity more private.

Each of those various solutions that have been prepared and are being experimented with are looking at a different piece of it.

So you're right. This does not make everything private. What it does is address a certain aspect of the privacy of your communication with he DNS. But it doesn't fix everything.

Now, the other part of what you said is what we cover in the paper. In the SSAC paper, we look at the different points of view—different perspectives, as we call them. Depending on whether you are a government or an enterprise or a service provider or an activist or parent, you may have different reasons you want things to be or not to be private and different people you trust and different parties you trust and don't trust. So the paper tries to look at that, and that is absolutely a significant issue with this: one person's protective filtering to keep pornography from our children is another person's censorship, and the perspectives will differ among different points of view.

So do read the SSAC paper when it comes out, which is imminent. I think you might find it interesting.

PAUL HOFFMAN: Alan, I want to give sort of the same answer that Barry just gave but from a very different perspective, which is you said, "Well, really this comes down to, who do you trust?" That's a nice, simple question, but I'm going to ask a counter-question, which is sort of where Barry went, which is, when you ask that question, who the heck is "you"? There are many "yous" in this discussion. One is the person who's sitting at the computer. The other is the person who wrote the software that that person chose to use. In this case, this is why the folks at Mozilla feel like they have a right to make these decisions for you. Another is the Internet service provider, or whoever is your IT department or whatever, because they know that, in fact, you don't think this about this enough and they do. In fact, if you work in a company, you are paying those IT people. You the employee in some sense are paying those IT people to think about something that you are not. To go back to one of the earlier questions, the you could be your local government for some value of local, where they feel like, "There are too many yous in this question. We are the ultimate you. We are going to decide who to trust."

So really, for the "who do you trust?", the answer is so complicated, not because trust is complicated but it's hard to figure out who's asking the question.

HOLLY RAICHE: Thanks, Barry. I think we've got 30 more minutes, but we haven't had time for, Barry, you to talk about the SSAC paper. When is that coming? I went looking for that yesterday and, from what Rod said, it's not on the website yet. Could you talk to the paper for about 10 or 15 minutes? Then we can go back to questions and the three hands that are up now. Thanks.

BARRY LEIBA: I can do that. The first question there is, when is the paper coming out? As I said, we are just getting the last edits done of it before it's released. So I expect it any time now.

Let me talk a little bit about what to expect in the paper. The paper does a bit of explanation on the order of what Paul did of what DNS-over-HTTPs and DNS-over-TLS are but doesn't spend a lot of time on that. It's mostly looking at the effects of those protocols on and the perspectives of different groups of stakeholders. We talk about parents, enterprise, network managers, dissidents, protestors, and Internet service providers—the different parties that may use or implement these protocols or that implement resolvers and are affected by the changes in behavior of applications. So that's the focus of the paper: looking at the different perspectives of different types of stakeholders. So we're looking at how the applications making their own resolver choices affect different parties and what implications arise from those decisions—the implications on the namespace that's due to changing the points of control along here and having stub resolution moving into applications, and different applications making different choices about which resolver to use.

I'll add that one thing that Paul didn't talk about, I think, in his part of the explanation is that, traditionally, when you used a stub that was built into your operating system, that stub would always use the same resolver. So all the applications that went through it would use the same recursive resolver. Now each application is going to make its own choice or has the opportunity to make its own choice about which recursive resolver to use. Different applications may make different choices. In theory, any recursive resolver may return the same information. In practice, that's not always going to be true. So the paper will talk about how that affects things and how splitting it affects things.

One thing that the paper does not do is say, "This is the right way to do things, and this is the wrong way to do things," because we looked at this and, as we discussed it, we said, "These issues are nuanced." As I said, different perspectives have different views of it. Different people in SSAC had different views of it. We're trying to avoid saying, "This is right, and this is wrong." It's, "These are the tradeoffs. These are the nuances. These are the different effects that different choices have." So we've avoided the kinds of statements like, "More privacy is always better," or, "More encryption is always better," because of those tradeoffs. So our recommendations are sparse in there. We're not really making recommendations in the SSAC paper, as we often do. In this case, we're laying out the issues, laying out the checks and balances, laying out the different perspectives, and hoping that this helps readers understand more where people are coming from when they discuss this.

I think that's all I wanted to say about that. Suzanne Woolf is my Co-Chair of the work party that developed this paper. Suzanne, do you have anything to add?

SUZANNE WOOLF: Let's see if I have audio.

BARRY LEIBA: You have audio.

SUZANNE WOOLF: Awesome. Configuring this strange docking station. No, I was configuring a docking station in a borrowed office, so I wasn't sure whether audio was going to work.

I think, just to pick up on what Barry said—I've been chatting back and forth in the chatroom also—there's an awful lot of complexity here. We tried very hard when we first started writing the paper to reduce that to simple recommendations or simple principles. It turns out to be quite difficult because the diversity of interests here is fairly complicated and fairly hard to trace. I think, as we've been talking in the chat, one of the things we're converging on that also came out of the SSAC effort was that any set of solutions, either technical or policy, that didn't take into account that complexity of interests was probably not going to ultimately be successful.

One of the things we've seen with the deployment of DoH and DoT is that, as … With a great deal of these technologies, first you invent the technology, then you have software that implements it, and then you have to configure it. You have to say which server you want to use. You have to say which policies you want to apply. A lot of what's happening is that people are rolling out these technologies, looking at how they work, looking at what interests are impacted, and refining their models of how they should be configured and deployed. So having users paying attention and, frankly, having proxies for users who are slightly better informed and pay more attention—this is what the SSAC paper is aimed at and what I believe the OCTO paper is also aimed at ... The more informed base we have among people who pay attention to things like how DNSSEC actually works, the better it's going to be as users and other proxies for users to try to find their way through what's the best way to manage the complexity of choices here.

HOLLY RAICHE: Thank you, Suzanne, and thank you, Barry. Just to note we're not close to running out of time, but there are lots of questions. Remember, all these sections will be recorded. There will be transcripts. There will be recordings. The chat has been really informative. That will also be saved. So, if we run out of time, just remember you can go back and look at the slides and you can listen to the session again.

Before I got onto the hands up, is there anything further that either, Paul, you are Barry want to say? Or can we go onto the questions?

PAUL HOFFMAN: Riffing off of what Suzanne just said, lots of people have a lot of things that they've been thinking about on this. I don't think that that's actually worth me taking my time to tell you more. I'd say let's answer more questions.

HOLLY RAICHE: Okay. Let's go ahead. Judith, you've been waiting a while. Thank you very much.

JUDITH HELLERSTEIN: Thanks so much for this great session. It's a bit technical in many areas, but it also, I think, answers some of the questions that I might have.

In Firefox, when you go to certain places in the world, certain sites that work normally don't work anymore or they give you false errors saying there's a security issue or it's not a security issue, but then they work fine on another browser. So is this the problem that Firefox with the DoH has? Or is something else that's being lost in the conversation that the browser is having with the servers or something like that? Maybe you can expand on that because I think it was a little bit too technical. But I'd like to understand it more. Thanks.

PAUL HOFFMAN: I'll take this one. It is technical, but it's mostly not technical. The answer to your question is, no, it's not directly related to DoH and things like that coming from Firefox, but it is un-technically related in the sense that, going back to Jonathan's question of, "Well, who do you trust?" when a browser is checking a website for, "Is this trustable?" they have to have some rules that they internally that they have built up based on their own rules, based on what they've seen you doing, and such like that. Different browsers build up those rules differently, and they build them up differently at different times. So, when you hit a website issue that seems to be location- or time-dependent, it is really a trust issue from the browser. We aren't seeing those currently with encrypted DNS but, when we have this discussion another five or ten years from now, I think the answer will be, "Oh, yeah. Those are sort of related."

JUDITH HELLERSTEIN: But why would the same website, if you are going from outside the U.S., be different than inside the U.S.?

PAUL HOFFMAN: Again, the browser is making choices for you that are not technical choices. They are trust choices. So they might go on a different path. They may say, "Oh, I know where you are, and where you are has some legal restriction on what I can show you here." There's a million of these non-technical decisions that are getting made behind your back on these things in the web world. Now that we are having encrypted DNS and that the operating systems and the browsers are going to have to decide which of the encrypted DNS recursive resolvers to trust, those same kinds of questions will be coming up again, starting in the future.

JUDITH HELLERSTEIN: Thanks so much.

HOLLY RAICHE: Thanks very much. The next question is from [Loretta]. Could you unmute yourself?

Apparently not. Can we go to [Shiva] then, please? [Shiva], do you have a question?

[SHIVA]: [Just a minute]. Can you hear me? Hello?

HOLLY RAICHE: Yes. Go ahead, please.

[SHIVA]: I understand that the discussion is about authoritative DNS servers and [inaudible] the entire DNS queries and traffic, but my question is, does the registrant have a role in all this? From what I understand, in the web-hosting interface, a registrant has to enable HTTPS and, if he does not enable HTTPS, the domain still gets served as on the HTTP protocol. And HTTPs is not forced. So my question was about the role of the registry and the [role's] intention to make the security measures. Is there a role for the registries? Is there a role for the registrants? I'm technically not so well-educated, but I hope I have made a question.

PAUL HOFFMAN: I can answer your question simply by saying, no, this is not related to that. Your question is reasonable to the larger picture of, how does the

DNS help increase privacy and such? But the registrant and the registrar are outside of the discussion we are having now about the queries that go [to] existing domain names.

HOLLY RAICHE: Thank you.

[SHIVA]: Thank you.

HOLLY RAICHE: We also have a question from Thomas De Haan from the European Commission again. Go ahead, please.

THOMAS DE HAAN: Do you hear me?

HOLLY RAICHE: Yes.

THOMAS DE HAAN: My question is related to the DNS capability for filtering. Basically, we have been approached by the commission, by a lot of [inaudible] stakeholders, member states, law enforcement, and ISPs, which basically come down to the local DNS filtering capabilities, which are [inaudible] are now affected. So we have heard from several sites that one of the possibilities to mitigate this is to have local DoH resolving made possible. I wonder if somebody—one of the presenters—who know how the progress is in the achieving this because I think, if this would not lead to some agreed-upon guidelines for DoH deployment, then you would really face major problems. Thank you.

PAUL HOFFMAN: I'll take this, unless, Barry, you want to do this, since we both are hesitant to answer any questions that involve the word "law"?

BARRY LEIBA: Right. You go for it.

PAUL HOFFMAN: Okay. Why, thank you, Barry. I owe you one. Thomas, your question is extremely difficult. I don't mean that it's complicated. I mean that it's difficult to answer for the very reason that, as Barry said, the SSAC paper tried to look at these questions from many different perspectives, probably more than I had tried in the ICANN paper. Legal perspectives, as you are quite aware, have multiple localities. So, when you say, "When a user is in (and then you name a locality)," that's actually multiple localities. At least in Europe, as far as I understand, at any given moment, you're in a city or a municipality—you have nicer words for it than we do in U.S. language. You're possibly in a region of a country. You're in a country and you're also probably in Europe for changing value of the word "Europe." And there might be different restrictions in those.

So, to your question of, “How will the browsers deal with, or can they deal with it?” the answer is they’ll deal with it very poorly, which is exactly how they’re doing it now.

I take it as a hint that your question is that you would like them to deal with it better. Certainly, most people on this call would like them to deal with it better but have little to offer in ways of doing it. This has been a question that we’ve had literally since the beginning of the Internet: how do deal with these things.

With encrypted DNS, it gets a little bit more difficult, just in the same way it was when we started having encrypted HTTPS 25 years ago, in that, if one of those restrictions is, “I need to be able”—I, a government—“need to be able to see what’s going on in order to make some policy decisions,” you now can’t see what’s going on. Or, if one of them is, “You can only do things that authorize if you are within one of my regions,” you can’t see that it’s authorized.

So I’m sorry, but my answer is going to be fairly negative about this. But the good news is that, with anything that gets fixed with respect to this question in any part of the Internet, not just the DNS but also certainly web and voice, those same solutions can easily be transferred over to open DNS.

Barry, does that sound sort of like what you might have said if you were braver.

BARRY LEIBA: [Sorry]. Well, I'll add a little bit to it because Paul talked about a person being at the same time in different jurisdictions. I'll add that the complication is that the different pieces of everything you're using are in different jurisdictions. So you may be a German who's also in Europe, but you are using a network that's managed in the Netherlands and software that was written in France and was sold by a company in the U.K., and your storage is in a cloud server in the U.S. Who is responsible for what, and what laws are in effect, is just way beyond my paygrade.

HOLLY RAICHE: I like the answers which says how difficult it is. I've got—oh, Thomas, do you have a further question on that one?

PAUL HOFFMAN: Holly, before we go forwards, since you and I have not met, when I looked at your biography, it says that you teach law, and you teach law in a country that was not on Barry's long list that he just gave.

HOLLY RAICHE: Yeah.

PAUL HOFFMAN: Can you speak to this, please?

HOLLY RAICHE: Well, I teach Internet governance. I probably have given something close to what Barry said, which is, when you have a global Internet, which has got bits and pieces all over the place, law tends to attach itself to countries. We think of national laws. We think of state laws. We think of local council laws. Their interaction is problematic and, frankly, makes a very good business for us. When you make that global, it becomes far more difficult. So, the first thing I say to my students after trying to explain what the Internet is to them is it's really difficult to talk about global regulation when, in fact, the concepts we have of law tend to attach to a jurisdiction where there are processes in place to make those laws that in some way reflect the citizenry they cover.

So that's a very probably not adequate way of passing the buck. Did I do that properly?

PAUL HOFFMAN: You're the professor. Don't ask me.

HOLLY RAICHE: It's a very difficult question because there are overlapping jurisdictions. And it's far worse in the Internet. It really is. So you wind up almost quoting Lawrence Lessig, who says, "Code is law." And you leave it there.

I think that's me passing the buck. Isn't it?

Paul, you've left me.

PAUL HOFFMAN: I have not left you. You are still chairing this session. Since you are the one person who had I seen who was identified with law and government, I thought that that would be a good addition because folks like Barry and I and Suzanne and Rod are not from the law world. The early questions about encrypted DNS were first about, "Hey, what about my local policy decisions that I made, such as my ISP, who is doing this?" Then the second wave was about law, and they weren't questions. They were demands. They were, "We expect you to do this. Why are you not already doing it?" which is a perfectly reasonable thing for a government which is used to doing that to do. In fact, I just see in the chat that Patrik Fatlstrom, who is also an SSAC member, is talking about what Thomas said. Patrik has been on the pointy end of some of these legal questions for longer than I have.

So these kinds of questions are reasonable, but they aren't reasonable to expect answers from the techies.

HOLLY RAICHE: Quite honestly, listening to this, when people ask my legal questions, I'm going to say, "Well, actually, it's a really complex question. It's a complex questions because, at law, there are lots of laws and they're probably contradictory." I will have the knowledge to be able to explain why it's so difficult and why there are so many jurisdictions and why, in the end, it's necessary to understand both a legal framework but also a technical framework. It means the answers are not easy ones. One thing I've got note of in this session is the fact that the answers aren't easy.

From Barry, I'm hearing that in fact your paper nicely spells out the questions and at least highlights the fact that there are questions that have to be answered. And it's not an easy task. But, in fact, both of you have done brilliantly to explain why it's not an easy task.

THOMAS DE HAAN: Can I just chime in just for one remark?

HOLLY RAICHE: Of course.

THOMAS DE HAAN: Thank you. Of course, the legal realities are completely different than the Internet literal reality. But, then again, I think that at least local resolvers, as much as local is possible, [have] a lot of problems. So [inaudible] user, preferably [inaudible] [network where he is.] Thanks.

HOLLY RAICHE: Thank you. Now, I only see one more hand. Is it [Labretta]?

Not talking. Okay. Are there any further questions? Just a reminder to people, there's Paul's paper. Paul, you put the link to that paper in your chat, haven't you? If not—

PAUL HOFFMAN: I have, yes.

HOLLY RAICHE: Good. Rod, you have promised that we're going to be able to read the SSAC paper fairly soon.

ROD RASMUSSEN: Yes. As I said, the ink is still drying on the PDF, as it were, but SAC109 will be published. We traditionally send that to the Board first, and then we'll send it out. Again, my apologies for not getting this out ahead of time, but we had to get through our bit.

We also had prepared a presentation for this which largely overlapped with Paul's, but there's some different things in there. I don't have the link in front of me, but Andrew or somebody who has that link to that paper, if we have that, could paste it in the chat—not paper. The presentation. That might be useful as well.

HOLLY RAICHE: Yes, please.

[ANDREW]: I will do so.

ROD RASMUSSEN: Great.

HOLLY RAICHE: That's great because that means that, for people who are like me—still absorbing all of this—we can go back, read the paper and read the slides and the presentations.

We are really at the end of this session. I'd like to thank you both. Well, actually, I thank everybody for all of the stuff that's been in the chat. It's been really interesting. Thank you, Paul. You've made something that was very complex a lot less complex. Thank you, Barry. It's been terrific and also helped us understand stuff. And thank you, everybody.

With that, I'll just end this session with a big thank you and hope everyone goes back and rereads and reads this again to understand some of the complexity that maybe you didn't get the first time. I know I didn't. Thank you.

This meeting is now adjourned. Thank you.

UNIDENTIFIED MALE: Thank you, Holly.

[END OF TRANSCRIPTION]