MONTREAL – SSAC Public Meeting
Monday, November 4, 2019 – 15:15 to 16:15 EDT
ICANN66 | Montréal, Canada

UNIDENTIFIED MALE: This is the SSAC public meeting, 513D, Monday 4th November 2019, from 15:15-16:50.

ROD RASMUSSEN: As you're coming in, if you are not an SSAC member we invite you to sit at the table or in the front here, because of the way the room's set up. We want SSAC members in the back, and everybody who's come here to hear us in the front.

UNIDENTIFIED MALE: Okay. Any guests to the meeting, please sit in front.

UNIDENTIFIED FEMALE: Yes.

UNIDENTIFIED MALE: Oh, God.

ROD RASMUSSEN: Did we capture that, so we can use it as a meme, now? Alright, thank you everybody, and apologies for the way the room was configured,

here. If you didn't get the note, if you're an SSAC member … We still have plenty of seats, so I guess you can sit everywhere now. We wanted to make sure that the folks who aren't in SSAC, who wanted to at least be able to see the screen, had the seats in the front, and the SSAC members who've already seen these things can sit in the back.

Alright, I'm going to go ahead and get started here. Welcome to the SSAC Public Meeting, here at ICANN66. Good crowd. Thank you for those of you who I know are chagrined to be missing tech day at the moment. There are some interesting presentations that I was able to see just recently. We'll see if we can avoid that conflict in future ICANN meetings. There's always some sort of conflict. At least we're not at 08:00 in the morning on Thursday when two people show up, and either me or Julie or, I guess, Merike and a couple of others who have to be here.

Anyway, let's talk about what we've been up to. We've got several updates since we last presented in Marrakech. Let's get to it. If I could have the next slide? There we go, thank you. Is there anybody in the room who is not familiar with SSAC? Yes, not including the members. Alright. I just wanted to check that because we have some standard boilerplate, and I will quickly breeze through that if there are no newcomers in the room. I was hoping that some of the newcomers might come over. We addressed them yesterday.

Alright, good. You were shy before, okay. I won't just breeze through it, I could at least give you an idea. Here are the things we're going to talk about today. The standard overview. Feel free to sit in the front, please,

with the way the room's set up. The members have seen these slides before. At least, they'd better have. We have an environmental scan of the threat space. This is something new. We'll talk about that in-depth. Then, we have SSAC106, which was released since Marrakech. That is our comments on the RSSAC governance recommendations that they came out with. Then, we had a correspondence around root scaling as part of the public comment period, there. NCAP update, Name Collisions Analysis Project, and then other Work Parties that we have going. That's what we're going to talk about today. Next slide, please.

As of today … Please feel free to sit up in front as you're coming in. We have 38 members, board appointed. We have a responsibility around SSR issues to report to the ICANN Board, and to the community in general, on issues that are outstanding that we see as developing, as we get questions as far as things are being developed in the policy process, and as threats arise in the general Internet world, particularly when they affect the DNS and things within ICANN's remit.

We have a wide variety of expertise in our skillset, and we use that to try and be able to fill in and inform the community on issues from a broad perspective, and hopefully have people on hand that can address issues as they arise in various new ways over time. We have 106 publications, now, SSAC106 that came out. We of course have multiple other correspondences, but our main publications are 106. Next slide, please.

Publication process for our main papers. We have what we call Work Parties. Those are folks who are SSAC members who are interested in a

particular topic area, that have expertise there. We'll decide that that's a topic we want to do some work on. We'll do research. We have regular meetings and put together a product that is then presented to the full SSAC for a review. There's feedback, and that'll get reworked depending on what the feedback was. If it's approved, then we'll publish that. In some cases we done actually get to an agreement or consensus. A few times we haven't published something on a topic but typically we do.

That's the main process there. Those topics can come from board questions. They can come from security issues that have arisen. They can come from long-term things. We're going to talk about a new way we're approaching some of that prioritization. Then, of course, recommendations for the board go there, and there's actually a whole tracking system for them to work and interact with the recommendations from ourselves, GAC, ALAC, etc. That's tracked. Then, if there's some policy advice that may get pushed off somewhere else. If there are some recommendations that involve implementation issues that either Org or someone else may do, that may get shunted over there.

Then, in theory, everything gets figured out. Those recommendations are closed out, as to having being actioned in some form. That's the ideal process, there. It doesn't always work that way, but that's the ideal. Next slide, please.

I already mentioned what the two recent ones were, so we'll get into that in a second. There's some information there. We don't really have

that Facebook page updated. We should probably take that off if we're not doing anything about it, as I'm sitting here looking at it. We're looking forward to the new ICANN website so that we can have a new SSAC entry on that website, when that comes out. We do have information there.

In particular what's interesting there, as well, is if you are interested in joining the SSAC, you can go there to find some information about the SSAC, what things we may be looking for, and how to contact Cathy over here to get the application. A couple of our members are stepping down at the end of this year. They'll be recognized later this week.

Unfortunately, one of our members, Don Blumenthal, passed away last month. We are going to be down to about 35 members by the end of the year. We are definitely going to be looking for a few more faces. We definitely want to see about some diversity, both from a technological perspective, geographic perspective, and all the other standard diversity ideals that ICANN has. If you know somebody who's got some background in technical aspects of the DNS, or security, etc., and particularly if they're in countries where they have a different kind of Internet architecture than a lot of North Americans/Europeans are exposed to, we would really like to expand our knowledge base there. Next slide, please.

This is a list of all our current work, the Names Collision Analysis Project. We're pretty much finished up with the organizational review. In fact, we are going to be turning in the last of our homework in December on that, and finishing up the implementation of recommendations from

that. That's about to be off the list. Then, we have some continuous work we've been doing around our working processes. Our strategic and environmental scale, as we are going to be talking about in a minute, was one of those items that is delved into that, and I'll get into that in a minute.

We have a DoH/DoT, or DNS over HTTPS/DNS over TLS Work Party. We'll get a report out on that. We're just firing up an abuse Work Party that'll coincide with a lot of interest within the community, and something we've had on our radar for quite a while. Excuse me. As I mentioned, we are always thinking about what are new threats that are being developed out there, and how those might be impacting things. Our work this year is actually tied into taking a holistic look at that, and we'll get into that in a minute.

The DNSSEC and security workshop is on Wednesday? Yes, Wednesday in the afternoon. This typically had been in the morning, but because all the cross-community sessions were scheduled in the morning, that has moved to the afternoon. If you're a regular participant in that, you will note that we've typically had a lunch for the participants. That is not going to be possible this year because of the movement. There is, for those who show up and stay for the thing, an alternative post-event that will be occurring. I'll just leave that there. Is that right, Cathy? Okay, good. I don't want to get in trouble.

Then, we have our team. We have a Work Party in our representation, the ePDP. We'll have a read-out from that. Then, of course, I already mentioned we're looking for new members. We're trying to shift to

more of an annual membership process, where we can bring in all the candidates and take a look at what people may have, as far as skills, etc., and bring them together, so we can have a more regular process, rather than the ad hoc we've been doing over the years. Can I have the next slide, please?

Here's a few things that we may be jumping into that we're having discussions on one way or another, usually on our mailing list. From the top … We do not have report-outs on this, so any questions you may have on this, I would take now. These are some of the issues we've identified as potential Work Parties, or maybe tied to work we're doing in current Work Parties. Just from the top, that top one has to deal with the way that DS keys are managed between registries and registrars. There are some operational things that are kind of funky, just because of the way things are set up. There might be some ways to better streamline that. It may help with DNS, a second option.

The hijacking attacks, this was more of the allegedly state-sponsored stuff that happened last year, as the impetus for that. These things obviously have been occurring at the consumer level forever. Then, .internal is around potentially talking about a reserved namespace for doing interesting things as a way of having a safe place to play that doesn't cause new name collisions, potentially.

Resolverless DNS, that is around talking about in particular web browsers using an HTP code to be able to request resources, instead of going through the traditional DNS resolution, which is a trend we see potentially happening.

Then, the concentration of the DNS infrastructure. This is an operational perspective, and largely regarding resolver infrastructure. We're seeing that. I'm sure many of you are cognizant of the way people have been changing how DNS operations have traditionally worked to new methodologies. There's a bit of a concentration there, and there's some implications of that, that we may get into.

Those are the areas we're thinking about having some work done over the next year or two. Any questions about these before we move on to the current work that we're doing? Great, Next slide, please.

Alright. Talking about our scan. Next slide. Over the course of this year, we decided to take a look at the entire naming, addressing, Internet namespace, etc., threats, and do a holistic view of that, and really take a look at any of those four category areas there. Going out and doing some original research, and bringing in information about things that have been done by various outside experts, etc., and bringing those all together so we can actually get a handle on what people think the threats are that are out there, and get a better idea what we might want to prioritize in our future work products.

We've been picking work products based on what our members think are important, but then we thought maybe we'd do a more scientific way of doing this and try and align that with the strategic plan that ICANN has overall. It turns out the board has been doing some similar work, so we've been discussing that with the Board Technical Committee today. Then, do some ranking exercises, etc.

Could I have the next slide, please? Oh, is that it for that one? Okay. Well, go back one, first, just for one second. Just to give an idea where that's at, we did this environmental scan. That was largely done by our excellent staff. Then, we also had the opportunity to have a research fellow work on this, did a ton of really good, original research, and brought together a document that we now have internally, that we may publish after going through some more refinement. That's to be determined, but we really would like to. It's an excellent source document on all of these areas.

In our workshop in September we did a first-cut, basically a full day's worth of work on it, split over our workshop on a prioritization exercise, and came up with at least a first cut of things that we think are worthy of diving into a little bit more, and better defining. As I said, we are in consultation with the Board Technical Committee on how we might line that up with what ICANN Org and the board are thinking of, so that we're making sure we're aware of, and addressing in a prioritized order, things that are out there that are major issues that we haven't necessarily done any work on before. Also, identify the areas where we have done work but we may need to update that, in the light. As you might imagine, there's a lot of [piece] that we can do with that.

The other thing that we're going to try and do out of this as well, and we're already putting this in play, is identifying gaps in our own skillset of the members on the SSAC. If we've identified an area of risk where we either don't have any members or enough members to do justice to a Work Party, we'll work on recruiting people in that space to help us

out. Any questions on that before I move on? Okay, next slide. I'm going to turn the mic over to Russ Mundy to talk about SSAC106.

RUSS MUNDY: Good afternoon, this is Russ. Next slide. Cathy's running the slides. The evolving governance of the root server system is an activity that is a direct result of work that came out of the RSSAC. The RSSAC is, in a way, a sister body to the SSAC, that deals with the Root Server System. We, as the SSAC, had submitted comments on the documents that came from the RSSAC, that the board sent out for public comment. We were quite supportive in the need to do this, and we had four specific recommendations to try to help the governance Working Group get its work done and underway. Cathy?

First recommendation is that the SSAC have a voting membership seat on the GWG. In fact, as far as I know, the recommendations are all being positively received. The recommendation 2 is that SSAC should not be given any operational roles in any standing committees, because it's stated in the document that basically SSAC is not an operational entity and shouldn't be involved in direct operational things. Recommendation 3 is that the decisions coming out of the GWG should be made on the basis of consensus, if at all possible. If voting is needed, then that would only be if consensus cannot be achieved by the GWG.

Recommendation 4 was that an ongoing oversight and review process be put in place to ensure that the RSS is meeting the commitments and remains responsive to the needs and changing evolution. That's the

specifics of what we had. Does anyone have any questions? I see we have some RSSAC members here in the room. Please, don't hesitate to raise questions just because you may be an RSSAC member. Anybody have any questions at all? No? Okay, thanks.

ROD RASMUSSEN: Thank you, Russ. Alright, moving on to our next bit. Next slide, please. This is SSAC Correspondence 2019-07. It sounds so sexy. We prepared this because there were more communications around a subsequent round of gTLDs. The gist of what we said in that, it reiterated much of what we said in SSAC100, which was an official response that we'd get done earlier. Late last year I think SSAC100 came out, if I remember right. We reiterated our advice there. We did add a bit new to our content around that. There was a strong feeling amongst the members around saying something about the phenomenon of having … Some of the new TLDs have really high concentrations of abuse, especially when you take a look at DAAR statistics, and the like. The root cause of that should really be taken a look at.

As I mentioned, we've got an abuse Work Party that we've just fired up. Next slide, please. Those publications that this refers to are listed there, SSAC100 and SSAC103. Then, a couple refer to these other ones. These are more around the root scaling issues that have been around for quite a while that we've commented on before the first round of expansion. To be clear, those publications and those documents we put out before are still relevant. We have the NCAP project going on to address concerns around collisions. There were some other things that we had

recommendations on. Those really haven't changed. Any questions on this? Okay, not seeing any. Let's move onto the next bit. It's snowing. I see a polar bear.

KATHY SCHNITT: Well, we've got nothing.

UNIDENTIFIED MALE: It looks like smoke to me.

ROD RASMUSSEN: Please bear with us while we go through these technical difficulties. The next bit is going to be … Actually, I think Steve, you're covering that because Jim's not here, around name collisions. There you go, it's back up. Steve.

STEVE SHENG: Thanks, Steve Sheng, ICANN Org staff, in support of the SSAC. The NCAP co-chair is unfortunately unable to make this meeting, due to a conflict. I'm standing in to give an update. Just a quick refresh, the Name Collision Analysis Project, the genesis of that was the board tasked the SSAC to conduct studies and present data analysis and points of view, and provide advice on two issues. One, specific advice regarding the .home, .corp, .mail. These are three strings that are indefinitely deferred in the last 2012 round. Also, provide general advice regarding name collision going forward. In the board request, the board asked the

studies to be conducted in a thorough and inclusive manner that includes other technical experts. Next slide, please.

The way the SSAC approached this is different from the traditional SSAC Work Party. In cooperation with the ICANN Organization this is the structure that is set up to run the projects. The key here is the project owner and the project director is the ICANN office of the CTO, with SSAC being the technical architect and the advisor. The customer is the ICANN Board, with the steering group. Thanks, next slide.

The SSAC envisioned to answer the board's questions. The board has a list of nine or 11 questions. Three studies are considered. We're currently in the phase one, the study one. That is a gap analysis that aims to properly define name collision, review and analyze past studies that work on name collision, and perform a gap analysis both on the literature as well as any gaps in the data. The intent at the end of study one is there's a decision point by the board on whether to continue additional studies.

Studies two and three are trying to respond to the other questions from the board, such as suggested developing a criteria for determining whether an undelegated string should be considered a collision string. That is a string that manifests name collisions. Suggested criteria for determining whether a collision string should or should not be delegated, what to do about the collision string, and how to remove the string from the list of collision strings. There are other requests about the mitigation options, and a final set of recommendations to the ICANN Board. Next slide. Did it stop? There are more slides.

ROD RASMUSSEN: Yes, there's one more slide.

STEVE SHENG: There are more slides, yes.

KATHY SCHNITT: It's delayed. It's very delayed.

UNIDENTIFIED MALE: It's slow. Somebody speed this Internet thing up.

STEVE SHENG: This is the chronological progress of the project. I just want to jump to the very end. In December this year, ICANN has selected a vendor to begin study one. This past Friday, the NCAP Working Group had a meeting where we went through the definition of name collision and also had a Q&A session with the contractor. A call of action for the members here is, please join the NCAP Discussion Group. Those are on the membership Wiki, where you can go and fill out the form and join the group.

I've been discussing with office of CTO. We may also be issuing an ICANN-wide announcement to call people of the progress of the project and refer people to join the NCAP discussion group. Watch out for that

announcement from ICANN. Next slide? I think that's it, that's the last one. Thanks. Any questions?

ROD RASMUSSEN: Let me put a point on the last bit that Steve mentioned. Now that we're actually really going with this project, I think there may be some people waiting to see it get off the ground. Now that we are, I really encourage you to think about joining, especially if you have some data or research that is relevant to what we're trying to do. Now's the time where that can actually have an impact as the contractor's got up and running, and is starting to sift through the various literature, research, etc., that's out there, that's been brought to the table so far. If you've been doing some interesting stuff where you're thinking there's an application towards name collisions, we really want to see it included into the work that we're doing now. Everything we do now helps us decide what to do going forward. Any questions on where we're at with NCAP? Okay.

We have various Work Parties that we have going. I'm going to have various folks that are doing these Work Parties come on up, and give a quick update. First one, I think, is the DoH/DoT Work Party. If we could move to the next slide? Either Barry or … Suzanne's name should be up there, but oh, well. Whichever of you two lost the arm wrestling wants to give us an update, I don't remember who. Is Barry here? Suzanne?

SUZANNE WOOLF: I guess I'm doing it.

ROD RASMUSSEN: I guess you're doing it. Sorry about that.

SUZANNE WOOLF: [inaudible], we'll see how this works. Barry and I have been chairing a Work Party in SSAC for … We started out by talking about some of the new transport protocols for DNS. There's DNS over HTTPS, and DNS over TLS. DNS over HTTPS is getting most of the oxygen in the discussion about the implications of these technologies. We've ended up spending most of our time on DoH. Both of these technologies are ways of having DNS queries over other transports besides UDP, which is the standard DNS transport. It doesn't change the content of a query, or what the right answer is. There's no impact on the database of what's stored in the DNS, or what answer you should get to a query. These technologies do change how a query is transmitted to a resolver.

The primary change that they make is that they encrypt the query in the response and inflate the ideas to reduce the opportunity for third parties who have vantage points in a network to be able to see the contents of a DNS query or a DNS response. Quite frankly, the DNS queries and response is generated by your device or apps, and so on, are a big part of your presence on the Internet, and leak a lot of information about who you are and what you do. The idea with the encrypted DNS transport is to reduce the exposure of information that might be nobody's business but yours.

There's been a lot of controversy, though, about the encrypted DNS transport, for various reasons having to do with the way they change access to who has information about what you're doing in the DNS. It also changes some of the mechanics of where decisions are made about how an application will behave, with respect to the DNS. There are some subtleties that, frankly, one of the big challenges is figuring out how to explain how these things work, and what exactly they change and do not change about how DNS works.

The Work Party, we're writing an advisory to explain the technical changes. The target audience is the ICANN community and the larger Internet community. The report is discussing impacts on different Internet actors, end-users, network operators, governments, ISPs. The profile of any technical change looks different for each of the players.

Risk and benefit analysis of different deployment models, which is very hard to summarize briefly. Basically, there are different ways, not only of using these technologies but of building them into your apps, your devices, or your networks. There are pros and cons to each of them. One of the complexities here is that there are issues that arise of who decides what resolver you're making queries to, and what other implications arise from who does see your data.

What the encrypted DNS transports do is they change who can watch your data go by, but there are still two endpoints to any transaction. There's still the source of the query and the resolver that has to receive it and reply to it. When you change who those players are, how those

decisions are made, you change some of the relationships among the players. Next slide, please. Sure.

ANTHONY EDEN: Anthony from DNSimple. Before you go on, I just have a couple of quick questions for you. Will the report compare and contrast …? Maybe that's going to be covered, here. You know what, it looks like you're going to talk about it. I'm going to let you continue, and we'll get back to it.

SUZANNE WOOLF: Terrific, because that's actually one of the things we're struggling with, too. The preliminary findings, here. Most of the concerns that are coming up … Certainly among various technical circles, the deployment of these technologies has been a very big controversy. The primary change that's easy to identify is that your web browser is making a decision about where your DNS queries go, instead of the configuration of your device by your ISP. That is a change in how the DNS behaves for the average user. There's been a fair amount of layer-nine controversy and policy tussle about who is gaining and who is losing access to information, and what the implications of that are.

The thing is, most of the concerns arise not from the protocols themselves, which are pretty straight-forward, technically … There are open IETF standards about how these work, and they can be implemented in a pretty straightforward way. There are choices that

have to be made in the implementation, deployment and operation of the encrypted DNS protocols.

It turns out to be a fairly complicated situation. Frankly, a lot of what we've discovered in the course of the Work Party discussion is that it's actually quite difficult to characterize many of these things in general kinds of ways. We're working right now on making sure that the primary purpose of the advisory is actually to address what the ICANN community needs to know about these protocols and how they work. There's actually been an enormous amount published, some very good papers, discussions and analyses about how the encrypted DNS protocols work.

One of the things we're working on is a list of … Rather than write a tutorial ourselves, we're probably going to just say, "Here are some really good documents that are already out there." There are complexities, and being able to explain the issues in a way that will benefit this community and the wider technical community turns out to be challenging. We're trying to finish the document, make sure that it says something useful that SSAC can really add to the background that people have available, when the usual process is SSAC internal review, and then publish, and gather community feedback.

I think we're still actively in development, but the end is in sight. We're close to a full working draft. There's a lot of subtleties here, and we're trying to make sure that we're adding something useful to the dialogue. You had questions?

ANTHONY EDEN: Yes, thank you. You're right, this is a very complex subject. I think one of the things I'm interested to see whether it's going to be in this report is the differentiation between DoT, which sticks to a defined port for transporting DNS queries, albeit secure, and DoH, which hijacks an existing port for alternate purposes. I think that's one of the biggest points of conflict, here. Will the paper compare and contrast those two, and discuss about what SSAC feels is the risks with that particular model?

SUZANNE WOOLF: I think it's safe to say quite a lot our discussion has been around exactly that sort of question. I'm going to pick up on another piece of what you said. Calling what DoH does in that context "hi-jacking," by itself, is a judgment that if we're going to make, we're going to want to make sure we're justifying. To those who want to deploy it, being able to include DNS as part of HTTPS is a valuable property that increases the protection of the data in flight by ganging it together with other things you would also be doing. I'm trying really hard to be neutral, here. I can't speak to a consensus about what's in the document that doesn't exist yet. I think the question is a valuable one, and the way we find ourselves asking it shows very clearly why this is a really difficult paper to write.

ANTHONY EDEN: I absolutely agree that the word I choose in that case, I do intentionally, and I think a lot. You're in a hard position because you are trying to take a neutral approach, so I just say I appreciate the fact that there's a group looking at this in an objective fashion, and it's true that the technology underneath is fine. In any one of these cases, the protocols are well-defined. I think it's the use-cases that are proving challenging to operators at all different levels. Thanks for doing that.

DAN YORK: I think this is great. I'm glad that the SSAC is taking this on. I've had a number of people asking me specifically for this kind of paper. I'm thrilled that it's there. Thrilled somebody else is working on it. One question …

SUZANNE WOOLF: That's what we're here for, Dan.

DAN YORK: One question is, realizing that you still have to finish it, all that kind of thing, do you have any guesstimate, remotely, of a timeframe for when this will be targeted to come out?

ROD RASMUSSEN: I'm going to say "soon," for some definition of "soon." We have a lot already done. We've been at this for a few months, already. Just this past week, we've been wrestling with how far to go. Are there some

other things we want to say beyond just the description and the risks, and all that? I would say it's not going to be out next week. It's going to be out well before Cancún, so somewhere in that range. That's vaguely helpful.

DAN YORK: Thanks.

ROD RASMUSSEN: I'd love to have it out before the end of the year. That's what I'm shooting for, but we'll see. We just have a few more minutes, then we've got two more things I want to make sure we covered. Jaap, have you got something real quick?

JAAP AKKERHUIS: On my way to here, I've been asked by a lot of people, "Were drafts available?", but since I know the answer I just thought to ask [inaudible] of community. As far as I understand, we never do drafts. It's always the full report. I just want to make this out in the …

ROD RASMUSSEN: Yes, traditionally SSAC does not put out for public comments things. There's a little bit of stuff we've done because of NCAP, which is unusual. If we get a lot of feedback based on whatever we do release, we have done updates to prior things in the past. You can think of it, we'll do a lot a follow-on if there's a lot of feedback that we say would

require an update or would warrant it. Next slide, if we could. I've just got two more content ones, here. This is the ePDP team. I think they're in session right now. Tara, do you want to give just a quick update on where things are?

TARA WHALEN: Sure, sounds like a follow-up [inaudible]. Hi, I am one of the alternates listed up there. I am here instead of being over in the ePDP, so thank you for that. As SSAC members, we are of course one of the stakeholders who are working currently on phase 2 of the temp spec for gTLD registration data. We are trying to ensure, as it says there, that we are consistent with our SSAC advisories. We are representing this, the community of the Security and Stability folks in this process, as we make this policy, which will affect a lot of folks in this ecosystem. Is there another slide? There is, it's it.

Currently, it's on phase 2. That is the access to non-public data, for folks who haven't been following along. The first one was more about the collection and publication in the public database. Now it's, "Who are the parties who will be able to access the non-public components of the registrant data?", who those parties will be. There's an accreditation component. Only accredited parties would be able to access the system. There are some open questions as to how accreditation will work. There are some questions around the disclosure decisions that are being made. We have a preliminary model which will involve …

There are other parties in the registry and the registrar side, as well as ICANN itself. It becomes, how do the queries go through? Who makes the assessment as to whether a query is legitimate? And, whether the data will be disclosed, and who well that works in that process. How much of that can be automated, given the requirements for large amounts of data to be processed through this? What are all the privacy concerns of the folks whose information is being disclosed? And, how all of that is done responsibly.

Right now, in terms of timelines, if that's useful for folks, we are looking now, fingers crossed, to have a preliminary report out at the beginning of December. Bearing in mind, some questions have gone to the European Data Protection Board, which also would be coming back around that time. If we hit that timeline then it should be January … Sorry, this should be ahead of Cancún, but that's our optimism. That's where we sit right now. Thanks.

ROD RASMUSSEN: Thanks, Tara. I have one more slide, and that's on our abuse Work Party. Jeff Bedser will cover that.

JEFFREY BEDSER: Hi. I think Mark Twain once said, when reading his own obituary, "The rumors of my death are greatly exaggerated." The rumors of progress in this Work Party are greatly exaggerated because we haven't started yet. Recently formed is literally … I've got my list of names of people who've signed up. We haven't had a meeting yet. Any questions about

what we're going to do, what we're going to say, what the output is going to be, could be held off until at least the March meeting, please. Beyond that, we have a long list of topics surrounding abuse that has been going around the community. We are very excited to see the community talking about abuse and taking the issue as a very serious approach, even to the point of defining, "What is abuse?"

We have a lot of items on that Work Party. There's the potential that Work Party will end sometime around 2030. We're really hoping for some earlier deadlines than that. Wow, I really didn't mean that. Please, stay tuned. We plan on having some progress. We'd like to get some progress before the end of the year. I've actually suspended holidays for all my Work Party members. I didn't actually tell them they're on the Work Party yet, did I? Yes, we actually plan on having some advancements on many pieces of this program by the March meeting in Cancún. Is there something else that you want to cover on that, Rod?

ROD RASMUSSEN: Yes, let me so that. I'm sorry, Tara, I was giving you the "hi" sign to finish up, because I was thinking we were done at four. It turned out we started at 15:15, so we have until 16:15. If there are questions on the ePDP stuff, we could take those, too, which means we'll have 15 minutes for Q&A. I was trying to get us done so we could get done by 16:00. It was just like, "Oh, my God, we got more time."

JEFFREY BEDSER: I'm happy to take questions. I don't have any answers, yet.

ROD RASMUSSEN: Yes, there's that. I think there are a couple of things we can talk about. One of the things we want to try and concentrate on early is looking at a framework, or some sort of way we can add to the conversation around how to properly define abuse, and how that might be used in conjunction with the things that come under ICANN's purview, contracted parties' purviews, etc. Those are things that don't involve having to go out and get data, do studies, and things like that. Those are more discussions, and bringing those things together. That's one of the things I hope that we can get out of this, earlier.

Another thing that we want to be able to do is really reach out to – I don't remember if you've mentioned this or not – some of the folks in the industry who are really doing a lot in this space, seeing if we can really understand the things that are effective, best practices, etc., and bring them. If we want to talk to that a little bit?

JEFFREY BEDSER: Yes, the Work Party structure has determined to have non-SSAC members invited to participate, not just review, along the lines of potentially contracted parties that have abuse programs, that have tackled these different issues. We can do comparisons as Rod indicated between programs for effectiveness, best practices. Also, take truly business-based decisions into running anti-abuse programs that are not a cost-center but actually are an addition to the value of the clean

eco-system, and a clean business, and how we can demonstrate those two things together where it's not just a stand-alone issue.

ROD RASMUSSEN: We'll be looking to gather information from the community, both from the folks who are trying to deal with cleaning up the mess on the outside, and those who are dealing with cleaning up the mess on the inside, so to speak. You should be hearing more from us if you're involved in those kinds of operations. We want to talk to folks about the best things that they see, and approaches that are being done out there in the real world towards dealing with these things. We're going to try and bring this together, and get rid of the anecdotal evidence, and work with real data on how to do these things. That is the last we had on our prepared remarks, so to speak.

There is one more slide, which is, "What are things you want to talk about that we didn't talk about here?" I will also take questions on the last two topic areas, as well, or anything that we've brought up today. This is open to the rest of the community for things you'd like to bring up. I'll start right here.

ANTHONY EDEN: Just a quick question. Do you have anything that you're missing from the current Work Party in terms of skills or areas that you'd say, "Wow, if we could really get somebody that has this right now …" This seems like a good opportunity to advertise the need, if there's anything.

JEFFREY BEDSER: That's a great question, and since it's not fully formed yet, other than the fact that apparently I'm the Work Party leader, anybody with skills that could contribute to that topic, speak to me after and let's exchange information.

ROD RASMUSSEN: We have not. We'd need to sit down and take a look at our own interest level internally, then discuss where we think there may be gaps, before we do more outreach potentially, on that. Thanks for the question. Other questions? Hope you didn't break that computer.

UNIDENTIFIED MALE: I'm a bit curious. I know with DNSSEC that helps with the authenticity piece, but DNS over HTTPS/TLS, that helps with the encryption. I'm slightly curious. Why aren't those being pushed side by side, simultaneously? I think that would make a lot of sense.

RUSS MUNDY: Thanks for the question. It is one that actually has a very long history associated with it.

ROD RASMUSSEN: [I thought you were going] to say, "Very long answer."

RUSS MUNDY: The history really starts back in the 1990s when the need for making sure people that were using DNS were getting the correct answer to their requests. At that point in time there was, I would venture to say, no concern at any significant level about protecting the transit of the information back and forth. DNS information certainly, at that point, was viewed as totally open. Anybody could see it, and it was not a problem that confidentiality was needed. It took a lot of years to get the authenticity in place and working. When that was done, or a long ways along, people then started to realize, "There is a problem if others somewhere on the Internet can recognize or can see and observe, collect or change, the destinations and so forth."

What happened, as much as anything, it's really a matter of time and how the Internet has evolved, DNS being just one of many of the protocols that move and change over time. It was the perceived requirements for security at the time that the work was started that drove the decision to do these two completely separately. Warren, did you want to?

WARREN KUMARI: Yes, I'll add to that. Back in 2014 the IETF published RFC … What was it, 7258? I think 7258 … Which says that pervasive monitoring is an attack. There were a lot of disclosures that came out around that time, the obvious one being Snowden when people took a look back and said, "Well, this is potentially really bad. We need to deal with this pervasive monitoring thing." That kicked off a lot of work in the IETF to try and add privacy to all the protocols, which usually involves adding

encryption. Part of that work ended up being the DoT, DNS over TLS, work.

After that was done – and full disclosure, I was the chair of the Working Group – we realized that that does a lot of good things, but it's relatively easy for somebody to just block the DNS over TLS stuff. If you're in a country and you really want to reach the Internet, you then have no choice but to just turn off DNS over TLS, start using normal DNS and then the [census] wins or you no longer get the privacy protections. That's what ended up kicking off some of the DoH work. Patrik looks like he's going to add to this.

Another big use case, which some people view as being more important, some people as less, is it would be useful in some situations for web applications to be able to do their own DNS lookups. From JavaScript, there's no way to be able to do a DNS lookup. One of the other use-cases was to be able to leverage that sort of thing. Where you fall between the two use-cases depends on your privacy paranoia versus shiny new web app side.

PATRIK FÄLTSTRÖM: Another thing I just would like to remind people about is that DNSSEC is really about signing and being able to validate the authenticity of a response of all resource record set that has moved all the way from the server, wherever the resource record set has been signed to whoever is validating that, regardless of how many hops in a potential chain of TLS, secure DNS hops that data has passed. These two mechanics

actually do solve two very different kind of problems. One cannot replace the other.

BRAJESH JAIN: I just wanted to raise … There are certain IP addresses which are not in the registry system. How does a DNS deal with such IP addresses if they surface and then they disappear? I'm told that there are more than about 15% of 4 billion IP addresses. Some are legacy, not in the system. There is no one who owns them. There are certain IP addresses whose ownership is not known. As a security [ideal], IP addresses surface, do some transactions, and disappear. I do not know whether that becomes a sufficient topic to work upon. Thank you.

WARREN KUMARI: I guess I'll mumble for a while, and then you can come up and mumble for a while. There's two different sets of potential IP addresses which you might have heard fall into that class. Some of them are the legacy set of IP addresses. There were IP addresses which were handed out, often, by Jon Postel before the RIR system was set up. These are things where the addresses were given to a person, or sometimes and organization, and in some cases the records have been lost to time. Generally, people know who has that. Some of the RIRs will let you bring those IP addresses into the RIR system. That then requires you signing an agreement with the RIR saying that now they're managing them, which makes a lot of people very twitchy, especially people who have had addresses for a long time. [inaudible] starts costing money, etc.

There's another class of addresses where who owns them is somewhat unclear, and that's because IPV4 addresses now have a lot of value. Some organizations will have the addresses, and they will be the person on record that the RIR thinks they own them, but they will secretly lease them out to someone else. They're allowed to do this, but it gets hard to know who actually has the addresses at that time. There's potentially a bunch of technical things that could be done to make many of these better.

There's a new system, RPKI, which is basically a cryptographic way to prove who owns an address. This is starting to be deployed. Actually, it's being deployed fairly well in certain regions, especially in Europe and Asia. Very much most places other than the US, for stupid political reasons. That way when a router gets a route with IP addresses in, it can validate to a reasonably good extent whether it should accept those or not. I'm not sure if that actually answered your question, or if it answered a completely different question.

BRAJESH JAIN: No, thank you.

MERIKE KAEO: I have a comment, also. The ASO is meeting here, and so I would recommend taking that particular question regarding the IP addresses to that community.

UNIDENTIFIED MALE: Oh, we've got a question.

KATHY SCHNITT: Yes. This question is from John McCormick. "Will SSAC be examining the link between discounted or nearly free registrations and DNS abuse, and web spam, and providing recommendations to ICANN?"

ROD RASMUSSEN: I think that that's a topic space that is going to be at least looked at by the Abuse Work Party. I will not speculate at this point as to what conclusions we may reach, and then, of course, based on that, what recommendations we may have. There's a lot of research. As I said, we want to get real data to drive any recommendations or findings first, and then any recommendations we would be to make. There are certainly plenty of anecdotal evidence out there for such a concern. I think that that's been published in many places. However, we're going to take a fresh look at this and try and, as we're doing with the DoH Work Party, etc., try and be very scientific about it and come up with good data-based answers. You had a question over here?

UNIDENTIFIED MALE: I just wanted to ask, you are looking into the implications of DNS over HTTP and DLS, but have you considered a work item DNS over blockchain since there are talks of some gTLDs wanting to implement that for their customers?

ROD RASMUSSEN: DNS over blockchain?

UNIDENTIFIED MALE: I said it like that for a purpose.

ROD RASMUSSEN: I'm trying to figure out how that works. The blockchain stuff has come up, but we haven't put those on top. Would anyone like to speak a little bit?

WARREN KUMARI: Yes. There are a couple of different ways that people are talking about. There's stuff like the .DNS people. I think they're technically called okTurtles. More recently, there has been the GNU namespace, or GNU naming system is what they're actually called, has been doing some interesting stuff with that.

UNIDENTIFIED MALE: Also Ethereum naming system.

ROD RASMUSSEN: Yes, there's also Ethereum naming system, which is actually, I think, better deployed. One of the potential issues with some of them, especially the GNU namespace one, is at one point they wanted to have their namespace within a specific area, that looked like a TLD. They wanted to have something like .GNU. They weren't able to get .GNU,

and so they're instead building a system where they're just going to have names scattered all over the namespace. If it happened to conflict with existing TLDs, the person who was presenting this in the last IETF meeting – sorry, the meeting before that – largely said, "Couldn't get .GNU. Oh well, we're going to just put these names all over the place. If it happens to conflict with, for example, example.com, whatever." That's just one of the blockchain things.

There are a number of different organizations working on ways where you can have a blockchain-type name. Some of these are so that you can prove that you own a name, but a lot more common use-case is they want censorship or to take down resistant names. There's this thing called Zooko's triangle which says that you can't have a name that is decentralized and unique and human-understandable. They're trying to design a system that can make many of those things happen.

The problem is if you have a distributed namespace it either means that you just can't do takedowns at all, which means that, if somebody comes up with a name and tries to use somethingsomething.cook, for example, you could run into issues. Child porn, things like that. If there's no way to do take-downs, you have some set of issues. As soon as you update the system so you can do take-downs, you now end up losing a lot of the decentralized nature of the system. You end up with something that looks kind of like ICANN again. You need somebody to manage the namespace.

There are a bunch of different potential issues with these. It's fascinating work. It hasn't seemed to be very widely deployed yet. It's

possible that we'll turn around tomorrow and a whole bunch of people will start using it, and will become wildly popular, at which time I don't know what really happens with ICANN.

ROD RASMUSSEN: I just want to let folks know, we have now officially passed our time. I'm not sure if there's somebody in this room next or not, but I'm perfectly happy to take a couple more questions since there's so much interest right now. Wendy had one, unless you had a quick follow up to that?

WENDY SELTZER: Sorry. A very quick follow-up to that question, just to point out … Wendy Seltzer from W3C. W3C has a Decentralized Identifiers Working Group that is working on some standards for a DID URI scheme for identifiers that might not be used for DNS-like purposes at all, but for other kinds of identifiers that might be on blockchains. There is work at W3C you might want to take a look at.

WARREN KUMARI: Sorry, if I had seen you sitting there I would have just punted the question at you.

UNIDENTIFIED MALE: My question was [inaudible] DNS over blockchain on purpose because I think that we should do a study on implications before it becomes a one-click option in browsers and widely offered by others. I think we

should do some things in advance and have opinions and known implications beforehand. I am a little bit shocked with the late study of DNS over HTTP and DNS over TLS.

ROD RASMUSSEN: Thank you. Thank you all for coming today. We really appreciate the good crowd. We will see you again, at least many of you, in Cancún. Please, if you can, the DNSSEC and Security workshop, remember that is on Wednesday afternoon. Thanks.

[END OF TRANSCRIPTION]