ICANN69 | Virtual Annual General – SSAC Public Meeting
Tuesday, October 20, 2020 – 14:30 to 16:00 CEST

KATHY SCHNITT: Hello, and welcome to the SSAC public meeting. My name is Kathy Schnitt, and I am the remote participation manager for this session.

Please note that this session is being recorded and follows the ICANN expected standards of behavior. During this session, questions or comments will only be read aloud if submitted within the Q&A pod. We will read questions and comments aloud during the time set by the chair or moderator of this session. If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, you'll be given permission to unmute your microphone. Kindly unmute your microphone at this time to speak.

For all participants in this session, you may make comments in the chat. To do so, please use the dropdown menu in the chat pod and select "Respond to All Panelists and Attendees." This will allow everyone to view your comment. Please note that the private chats are only possibly among panelists in the Zoom webinar format. Any message sent by a panelist or a standard attendee to another standard attendee will also be seen by the session's host, co-host, and other panelists.

With that, I'm happy to hand the floor over to Rod Rasmussen, SSAC Chair.

ROD RASMUSSEN: Thank you, Kathy. Welcome, everyone, to another virtual SSAC public meeting. I'm glad to see lots of folks on board, not just SSAC. We have several things to talk about today, so we'll get into our schedule here.

Next slide, please, Kathy. There we go. For those of you who aren't familiar with the SSAC, we have a little bit of standard material. I'll go over that first, just as an intro. Then we'll talk about the most recent SSAC advisory. We put out SSAC113 on private use TLDs. Then we have a couple of items in response to public comments that we've recently had, and then updates on the NCAP (Name Collision Analysis Project) and other work parties that are ongoing within the SSAC. Then we'll talk a bit about the SSAC desire to add to our membership. [Dubee Smite] has some details on that that'll hopefully pique your interest or thoughts on somebody else you might think might be interested in joining the SSAC. So we'll finish up with that, and then of course your questions at the end—other security topics. We have 90 minutes in this session, so hopefully we'll have some time for community questions as well. Then we'll also take questions on each section as we go through if there are any in the audience. Kathy already mentioned the Q&A session there, so we'll try and answer things as we go.

Let's get into this session. Next slide, please. SSAC stands for the Security and Stability Advisory Committee. Hopefully, you already knew that, but that is the official title. We currently have about 35 people. We have a selection process which is based on a membership application process, which we'll talk about a bit, as I mentioned, at the end. But then, once the members have been chosen by the SSAC, that

is actually approved by the Board itself. So there is a formal Board advisory role here.

So we try to stay on top of security and stability and resiliency issues—SSR is the acronym we have for that—within the ICANN community, particularly issues that crop up, and then, overall, security items and areas that may affect the overall DNS and the Internet operations and affect the DNS. We draw from people with a wide range of experience—the whole list there—and try and bring different perspectives together so that we can cover different topics with our small-but-high-expertise group of folks to make sure we have coverage in those areas. 113 publications since the funding in 2002.

Next slide, Kathy. I think I have my connection back now. Yeah, there we go. Thank you. My apologies. My laptop decided to take a timeout. We have a process for our work that includes a formal work party within the SSAC. We create a work party for most of our topics. Some of the things we do within our Administration Committee or run past membership without forming a work party. Most of the time, we have a work party that forms based on interest and background in a topic area that then takes a look at a particular issue and works [inaudible] to a [inaudible] on the topic and deliberates on findings and conclusions and any recommendations that research and that work may bring forward. That's put together by the work party, which comes to a consensus on that. It's then shared within the full SSAC for a review by the full SSAC for further thoughts and refinement. There's an iterative process where that may go back and forth a bit between the full SSAC and the work party. Then, eventually, something comes

out of that where we have a full consensus document that is published. We may, if there are some differing opinions on the final product, we'll include those as well, which happens from time to time. That will be published. If there are formal recommendations, those will be shared out. Those are typically to the Board, and then there's a process where the Board takes in our recommendations, does an understanding, and then a little bit of back and forth between us/SSAC and the Board to make sure that's understood. Then any particular work that may come out of that is then tracked through ICANN Org or wherever it gets farmed out to.

So that's the process that we follow on doing that. It provides for a pretty robust examination of the issues involved and also of tracking through and seeing how that eventually may end up in some sort of implementation, which we've done a better and better job doing over the years. Actually, I think we do have some updates on where we are on that as part of the presentation today.

Next slide, please. We have two publications—we'll talk about both of these—that have been pushed out since the last ICANN meeting. The [O's probably use] TLDs, and our minority statement is part of the EPDP process within the GNSO world that I think everybody in the ICANN area, I believe, is familiar with, hopefully, at this point. So those are the two, and there's some information there around some of our outreach, which we'll talk about a little bit as well.

Next slide. First up is the most recent publication we put out a couple weeks ago. Warren, I see you're on the call, so I'm going to turn that

over to you. We've got a couple of slides on that if you would like to go through that.

And you are on mute, Warren.

WARREN KUMARI: Yeah. There we go. Sorry. The unmute button was grayed out for me. Hi, everyone. As Rod said, I'm Warren Kumari.

If we can have the next slide. Thank you. As always, you should actually read the SSAC advisory. This is just going to be an overview and some background. The advisory is relatively short and is actually a really easy read, so please read it.

So what actually is this all about? SSAC has previously advised that, if people need a private use space—something where they can number into internal devices or name internal devices—the best thing to do is to register a public name, like example.com, and then just create a subdomain of that—something like internal.example.com. Then, if you need to name things, you can call things printer.internal.example.com or fileserver.internal.example.com.

However, as we've all seen from things like the name collisions work, enterprises and device vendors don't always do this. The really well-known examples of this are things like .home, .corp, and .mail. But another really good example of this is .belkin. This is a case where a company needed a name which was for private use and just chose their company name and started using it as though it were a TLD. The other example is .home, which is where .home is being used by a

home router vendor, primarily in the U.K. They needed a private use space, so they just chose a string at random and started using that.

Unfortunately, the DNS does not have a way to prevent this. This means that, if people just start using names, you end up with name collisions. This obviously is not good for the namespace. It's not good for predictability, etc.

So what SAC113 does is it recommends that the Board choose a name and reserve the string using a bunch of criteria which we provide—basically, take a string, set it aside, and say, if you want to do this internal namespace stuff, this is a sanctioned place where you can do it. SAC113 specifically does not recommend a specific string, but what it does do is it sets aside some criteria for what the string should be like.

Next slide. Thank you. These are the criteria we believe the string has to have. It needs to be a valid DNS label. It cannot already be delegated in the root zone or already be in use. And it cannot be confusingly similar to another TLD that is already being used.

The final bullet point—this is where things get a little trickier—is it needs to be relatively short, memorable, and meaningful. The last of these is the trickiest to explain. As we said, there's no way to stop people from just taking a string and starting to use it. This means that whatever string is set aside or reserved for this use needs to be attractive enough that people will actually use it and not just continuing using a string that they would like to use instead. This means that the reason that people started using home, corp, and mail

and not supercalifragilisticexpialidocious, snickelfritz, or [inaudible] is because they wanted the string to be something that they could actually remember and use and meant something to them. So that's what we're trying to communicate in the "relatively, memorable, and meaningful." It needs to be attractive enough that people will be willing to use it and not just choose a different string because they can.

The SSAC believes that, if this is done, and if a string is reserved for this, not everybody will actually use it. As I've said a number of times, we can't force people to use it, but, if we create a sanctioned place where people can create internal names, some people will, and that will reduce the ad hoc usage of strings. This cuts down on the occurrence of name collisions in the future. It also provides greater predictability for network administrators and equipment vendors. So, when you see a string (whatever is chosen) and you see something ending in that, you will understand that it's supposed to be for internal use. If it shows up in diagnostic logs, etc., you will understand what it was supposed to be and that it has leaked.

As a nice side benefit from this, it will hopefully, over time, reduce the junk queries [that are pitched] to the nameservers. This last one is more just a nice side effect, not one of the primary goals.

I believe that that's the end of this. Rod, are we doing questions now, or are we doing them all at the end? Or how are we organizing that? Or, I guess, whoever is facilitating this.

KATHY SCHNITT: Rod? Oh, Rod's [inaudible].

JULIE HAMMER: Rod, I think we could take questions now, couldn't we? It's Julie speaking.

ROD RASMUSSEN: Yes, please. Is that feature locked in, Kathy? Because—

KATHY SCHNITT: Yeah. There's nothing I can do about that, Rod. I'll just try to be quicker with it.

ROD RASMUSSEN: Okay. Thanks. So, yes, questions, please.

WARREN KUMARI: Maybe there aren't any and I covered everything in full depth. Just for giggles, I have just pasted in a link to the ICANN DNS stats thing from the ICANN L-root server, which shows an occurrence of how often this sort of thing is happening if [I] just scroll down on the left and have a look at some of the names. You can get some idea of how often this is being used by things like home gateways. It also gives some sort of idea of some attractiveness-of-string-type ideas.

ROD RASMUSSEN: There are two questions in the Q&A.

WARREN KUMARI: I posted to the wrong place because I'm an idiot. There we go.

ROD RASMUSSEN: Warren, there are two questions in the Q&A. I don't know if you got the Q&A open.

WARREN KUMARI: I …Oh, there's the Q&A. Yeah, so two different organizations would be able to use the same private TLD. Currently, lots and lots and lots of organizations are using .corp, for example. This would be exactly the same sort of thing. My printer.internal and yourprinter.internal would be the same name but different devices. So the name would only be local to a specific—the correct term here is "domain," but that gets confusing—network or area. Hopefully, that answers that.

Jothan Frakes asks … Yeah, this is exactly the—well, conceptually—the equivalent of RFC 1918, but for names. So it sets aside a set of space which is intended to be used internally for this sort of thing.

There's also another one, which I'm not going to try and pronounce the name because I am sure I will butcher it. So SAC113 specifically does not recommend a name. It leaves that up to the ICANN Board and community. We suspect that there will need to be more than one of these, or there may need to be more than one of these, possibly because they're slightly different uses, but more likely because it's

going to be really hard to find a string that is memorable and meaningful in multiple languages. I'm not sure if that answers [both].

Vittorio is asking about SSAC's opinion on .qm and .zz. I don't know if I can speak to that. I don't think that the SSAC has a view on if that's short, meaningful, and memorable. I have a personal view, but—

ROD RASMUSSEN: I think we can agree it's short.

WARREN KUMARI: It is short. So I don't think I can give a better answer to Vittorio.

ROD RASMUSSEN: We did not come to any conclusion on that and any particular name. And the other question—

WARREN KUMARI: And, actually—yeah. Peter Koch has a question in the chat instead: "How would "meaningful" be defined in a global context?" That's why potentially there may be a need for more than one. If one has a look at the list of top NX domain queries which hit the root servers—the top bunch of these which are being used as names that people have just squatted on; I posted the link in the chat; I'll post it again—most of them—in fact, almost all of them; the top X—are sort of English words, but that obviously does not reflect the fact that there are many people who do not use English words. So this is going to be very tricky to

figure out what exactly the string should be. Luckily, that's more of a policy question than a technical one.

ROD RASMUSSEN: Okay. Just to answer a little more thoroughly the question about a list of common ones, there's no plan on creating a list of private use TLDs used by manufacturers. This is not the intent of this document, but you may see some of those in the list that Warren posted up in the chat that are leaking. The idea is that, if you were a device manufacturer, for example, and wanted to use something like that, instead of, for example, using your own company name or product name, you would use this reserved string for that particular usage and not add more random names to the quasi-DNS space, so to speak.

All right. Looks like we got all the questions, so thank you, Warren, for covering that.

I think we are going to move on to the next section now. The next was SAC112. Actually, we're going to talk about these two things. So let's move on to the next slide. I believe the next one was 112. Yes, it is. We have both of our EPDP representatives on. I don't remember which one of you volunteered to do this. Was that Ben or Tara? Tara. Okay. And you are enabled. I see you're unmuted, so I will pass it over to you.

TARA WHALEN: Thank you, Rod. As Rod said, I'm Tara Whalen, here with my able partner, Ben, and some other folks here on the call. We've been working hard as part of the EPDP process over the last approximately

year or so. We got around to the final report on Phase 2, which was the access to a non-public data component of this policy development process.

Where we landed: After all of that work, although we came to consensus on a number of the recommendations that came from this report, we did have some places of divergence and wound up in a place where were not able to endorse the final report in its current form from that we believe that there are places where we could come up with a better system under the limitations that are imposed by regulations, such as the GDPR. But, given our concerns around security and stability, we don't feel that the outcomes were able to address many of our concerns in that area.

On top of those recommendations, we also mentioned that there was not a commitment to finish some of the unaddressed charter items. So there were several items around legal and natural persons and accuracy of data that we didn't have any idea that these were going to be addressed and they had been left unaddressed for some time.

Now, we have had a recent meeting with the GNSO Council—I believe that was last week—in preparation for a lot of these meetings, and we believe there will be some progress on that front towards actually being able to address these items in the near future. I'll highlight four specific recommendations here, which are the issues that areas of particular concern to SSAC. I'll also note that the documents that we put out for these are SAC111 and SAC112. So this is the main commentary in 11, and our minority statement is in 112.

The four specific recommendations were primarily around response time and priority levels. There were a number of priority levels for response to particular types of requests. We have concerns around when there is some kind of an event happening where we require a quick response because we're trying to shut down some sort of a major attack, for example, where we need a rapid response when there's a request for information, where would like to see a shorter request period, where you get an SLA to get the information back. The priority levels don't really respond well. They're inconsistent with our needs, really, for being able to respond quickly. We would like a little bit more of a quick response. We don't think that was expressed well in the document, and that was really around 6 and 10, where we'd like to see a little bit of a better requirement around short and responsive times for particularly … We call them "time-sensitive" attacks—major attacks on things that are on critical infrastructure.

There was also concern in Recommendation 12 around the disclosure requirement around when a party has asked for information in this system, where there are concerns around the identify of a requester being disclosed when that is not actually required by a requirement of the GDPR, where, for example, the requester's identity … And instances where this could be some kind of … It would be difficult if a person was doing a request and if their identity were released when this was perhaps a criminal-type investigation, for example. You wouldn't want to have that person's identity revealed during the course of such an investigation.

And there's final recommendation around financial sustainability, where we felt that we would like to see more discussion around how the financial sustainability would be worked out in future … like how we would be able to figure out the ways in which this would work well, that there would be more of a consultation process and more information gathered in order for us to make some good determinations as to how the system would be sustained over the future. Those were our areas of primary concern.

We only had one slide, I believe, today. That was in brief. But we want to open this for questions before we go to the final report then?

Looks like everyone was following along, which is great. So thank you all for that, and thanks for the people who were working so diligently. With that, I guess I will hand off to the next phase of our presentation.

ROD RASMUSSEN: This is fun with the unmute controls. We'll keep me quieter. Before we jump on, any questions on this area? Because I don't want to short-shrift that. Oops.

Okay. I'll go ahead and cover SubPro's work next. As many of you have been following along, the Subsequent Procedures PDP put out a final draft for comment at the end of August. If you really are following along, you may have noted that the SSAC did not submit, at least as an organization, a comment.

Unfortunately, that time period overlapped with our annual workshop, which we did virtually this year for the first time ever, I

believe. I'm sure it was probably the first time the workshops were all in person or didn't exist before that. So we're not able to get things together in the timeframe.

But we did have some comments and thoughts in that area, and we actually met yesterday. It says we're a planning to meet with the PDP leadership. We actually met yesterday on the regular Subsequent Procedures PDP call—a 90-minute discussion with them on various areas that we had flagged either as a concern or areas to understand better or get some feedback both ways and have a conversation. I think it was a very productive call yesterday. [I] highlighted some issues. I know l saw lots of notes being taken. We'll likely be putting out some thoughts ourselves on that but we've at least had a good opportunity to interact with the SubPro team on some of the major issues, and even some of the minor ones, that we've identified at this point. Well, it hasn't gone through our formal process yet. We at least made them aware of what some of our thinking was in that process and where we stand today.

The two kinds of things that we talked about and concentrated on were the areas where we have previous advice and how that was handled in that draft final document. For the most part, those areas where covered. There were some areas that were considered and dealt with either differently than we might have indicated or what-have-you, and we had a good conversation about that. Going through that helped, I think, both the SubPro team and the SSAC members on that work party, understand a bit better how things ended up where they did, which is good.

The other part of it—I think this is an area more pressing or more a "Where do we go from here?" question for the full community—is about the areas that were not covered or gaps that weren't in the final report and in particular where the SubPro team have looked at an issue and determined that that was more of an issue for all TLDs or all gTLDs at least versus just new TLDs. One of the particular examples on that is a question around DNS abuse, which was pushed over to the SubPro to consider from SSAC advice, from CCT Review advice, and, I believe, GAC advice as well, if not even more. The SubPro team pointed out, rightly so, that abuse is an area that touches all TLDs and needs to be considered across the space, which it does. The question is, what's the approach, and how does dealing with that DNS abuse issue or any of the other issues that were flagged as needing to be looked at across all gTLDs … how do those fit in with the process of going into another round or set of rounds of new TLDs, and what are the things that need to get done by when in order to do that?

Then also one of the things I think the SSAC wants to point out is that sometimes there are some things you can try in namespace that doesn't have any names in it yet that'll be a lot easier to try there and see how that works then trying to retrofit (for a better word) a namespace that includes some new policy or technical provisions. If you recall, in the last round, DNSSEC was a requirement for all new TLDs, which was not necessarily a requirement for all TLDs at the time. That actually exposed several operational items that were people were able to learn from and then apply across the space because I believe [there's a] requirement to all TLDs at this point that took

contracting to do. It wasn't a policy thing, if I'm remembering that correctly. Somebody, I'm sure, will correct me if I'm wrong.

So those are the areas that we talked about. We have obviously the Name Collision Analysis Project, which we'll talk about here a little bit more, which fits into this as understanding of what criteria [there are] for names that may create issues which make them difficult to add to the root zone without some sort of mitigation or even to the point where the mitigation may not be sufficient. We are working through that right now. But that obviously ties into this work as well.

So those are the highlights of where our thinking is on that right now. So this is a work in progress still. We'll probably have some comments to pass along. As I said, having a conversation with the SubPro team, hopefully they will be able to accommodate at least some of our thinking there. I think a lot of our findings/recommendations, should we have them out of this, would be around these more meta issues that weren't addressed by the SubPro team because of scoping and where do we go from here on that.

So that's just a quick overview on that. Are there any questions or comments on that before we move on?

KATHY SCHNITT: Rod, just to note, your mute and unmute is fixed now.

ROD RASMUSSEN: Thank you.

KATHY SCHNITT: Thanks to our fabulous techs who found the issue.

ROD RASMUSSEN: I'm not seeing anybody tapping anything in the Q&A, so we will go ahead and move on. If somebody comes up with a question, it's fair game at the end. We'll go back and get it.

So let's move on to the next slide, please, Kathy. The Names Collision Analysis Project, which I just foreshadowed. Jim, I believe you were going to cover this one today?

JIM GALVIN: Yes. That's good. Thanks, Rod. Let me acknowledge my Co-Chairs: Patrik Faltstrom (also an SSAC member) and Matt Thomas, our community member who joined us so that we have a trio.

Next slide, please. This is just a quick look at who we are and what we are. Presumably, folks have been paying attention to the last couple years here. Just quickly, the Board had asked SSAC to conduct studies and put together an analysis and point of view about name collisions specific to home, corp, and mail, and, then of course, general advice going forward. So, since we are prepping to launch a new round of gTLDs, they wanted some advice on how to proceed when they get applications in the presence of name collisions.

We currently have 25 discussion group members. It's an open group. It's open in the sense that anybody can join, although you do have to

fill out a particular conflict-of-interest [form]—an ordinary ICANN process. Most people have such a thing on the ICANN wiki as part of the community. There are a few additional questions for joining the NCAP group that we do ask you to respond to. That gets posted up there with the NCAP project wiki.

I want to point all this out because we are always interested in more people. We are going to be getting ready to jump into our Study 2 here. Next slide, please. We'll be very interested in adding additional people. So, if you have an interest or you know someone in your group, especially if you are a registry operator or applicant or considering such a thing, we'd love to have your participation, as well as anyone who has got any data analysis skills with them to join us.

We did have Study 1 complete. We had a proposed Study 1 that was worked on at the beginning of December of last year through January/February of this year. Then we went to the usual ICANN comment process and such. So a final proposed report was put out for comment in May, and we delivered that report to the ICANN Board—the Board Technical Committee, in particular— at the end of June, completing Study 1. Study 1 was largely a bibliography of everything that we know or have learned about name collisions since 2012, when the last round launched.

In addition, there was a question there that was asked of the contractor to comment on their thoughts, given everything that they have seen[.] Having put together that bibliography, [are] Studies 2 and 3 ready be to kicked off and launch forward[?] Did the contractor have

any view about that? The key component of the response that the contractor had was that those studies, as designed by SSAC almost two years ago now did not really align with what we should do. It didn't seem that they were currently designed in a way that would be most helpful for us. And the NCAP Discussion Group agreed with that.

So we are currently in the process of revising Study 2, hopefully not in a very significant material way but sufficient so that the analysis that needs to happen so that SSAC can respond to the Board's questions and also provide specific advice regarding corp, home, and mail [that] will be relatively straightforward to jump to.

So that's where we are at the moment. We hope to have all of that done in the not-too-distant future. We're not really fitting to a timeline. We're going to do our work as efficiently and effectively as we can. But, hopefully, before the end of this year, we'll certainly have submitted a revised proposal to the BTC after it has gone through a discussion group review and an SSAC review. Then we'll just see where we are from there. Ideally, we'll enter a process in which we will be able to execute on Study 2 for now and 3 in the future, although not too long into the future.

That's it from me. Thanks. Any questions? Happy to take them or move it on. Back to you, Rod.

ROD RASMUSSEN: All right. Any questions on NCAP?

I'll give it just a minute here.

All right. Well, thank you, Jim. Much appreciated. We do have that study group, as Jim talked about. So I encourage you, if you are interested in this topic and have been participating, to think about that because we're going to get into some interesting stuff, I think, here, before too long.

All right. Let's move on to the next slide then. Current work parties. Alrighty. Let's see what we got up first.

Next slide. Oh, hey, it's a list. That's what we have first. I was trying to remember—no, no, no. Yeah, there we go—no. Back up. Back to the list. I'll just talk to the list really quick. We already talked about NCAP. One of the things that isn't on here that we do have a work party on the Subsequent Procedures' response. That's actually current work as well. Then we're going to have a couple slides on some topics here. There's some ongoing things that we do. Tracking SSAC advice to the Board—a little typo there. Sorry about that. We do have slides on most of those.

So now let's go ahead and go on to the next one, which is DNS abuse. Jeff Bedser, I'm going to turn that over to you to give a quick update on where we are there.

JEFF BEDSER: Thanks, Rod. Thank you for the tech crew; I got to unmute my own mic. Much appreciated.

So this is an upcoming paper that SSAC has been working on for quite some time. We have quite a larger charter to deal with DNS abuse. This

is probably one of the first efforts to come out of that charter—this current paper—and it just addresses the issues and that DNS abuse is expansive. It's not really an issue of, is it growing or shrinking? It's a function of a changing dynamic where the criminals and fraudsters are always looking for new ways to victimize people. This is about looking at new tactics to reduce it, not just through identifying it but also looking for ways to more quickly identify and more quickly react to it to reduce the number of people victimized on the correlation of: the shorter time it lives, the less people can be victimized, and thus the less losses.

We are in the final stages of this work party going to a paper. We're in SSAC review from the work party to SSAC. The hope would be to have it soon enough to share with the other parties. We discussed this a bit earlier in the DNS abuse plenary discussion. It feels like a day ago now, but it was just a couple of hours ago—this morning for me.

The gist of the effort is that we're taking these key points forward, encouraging the adoption of a standard definition of abuses.

Determine the appropriate points/primary points of responsibility for abuse resolution. At what point does a reporter of abuse go into the ecosystem to report it? That, again, is keeping in mind that the parties to address abuse or not purely contracted parties with ICANN. They can be entities that are outside of the ICANN ecosystem, such as hosting companies and CDNs, and mail providers, etc.

Identify best practices for the deployment of evidentiary standards. So, when somebody is reporting, for example, a command-and-

control botnet, the standard on how you report to any entity that's going to try and react to it is the same. This is not the evidentiary standards to meet a criminal justice issue within a different nation. This is about a standard definition and its standard evidence to present and provide on a DNS abuse domain.

Standardized escalation paths for abuse resolution. So, if you go to one party and they do not respond or will not respond or you cannot identify who that party is to take the action, where do you go next to get it escalated. Determining reasonable timeframes for abuse reporting. So, as there's an escalation path between parties and as there's a reality on validation of information that's been presented evidentiarily, what is a reasonable timeframe for the abuse to live until it's taken down based on responsiveness and timeliness?

Recommendation on the development of notifier programs that'll help expedite the process through channels of who you approach and how you approach them and why you approach them based on where it fits in the ecosystem, what evidence is presented, etc. As you can see, there's a progression happening here.

Creating a mechanism for the availability of contact information for abuse mitigation. Simply put, that's not registrant data. This is about the entire ecosystem having a methodology and a standardization around, how do you find out who to contact for a certain type of abuse and where to find that information?

Lastly, creating a mechanism to ensure reasonable quality of that contact information—timely, updated—to assist in quick reporting of abuse.

So this is a document that is really about trying to move the ball forward, reducing abuse not through stopping it from happening and not reducing abuse through faster identification necessarily. It's about, once it has been reported, how can we expedite the process in an environment where we're all cooperative parties trying to resolve an issue from across the ecosystem.

I think that covers it, Rod. Back to you, unless there are any questions.

ROD RASMUSSEN: Any questions or comments or input on our work here so far?

I'll give folks a minute. If you have a question … I'll let you take care of any questions if you got that up, Jeff.

JEFF BEDSER: Sure. Heather [inaudible] asked, "Can you so more as to why SSAC will not provide a formal definition of abuse?"

We're not rewriting the definitions of abuse. What we're basically taking is industry standards that have been put forth so far. In our paper, we do reference the Internet [&] Jurisdiction recent paper as well as the work from the framework against DNS abuse to start their process on definitions. My belief is that the definitions are continuing to grow, as there's always new types of abuse coming out or new

means or methods to the same abuse. So the definitions are not as important as reacting to the abuse itself.

I think that was the only question, Rod, so, I guess, back to you.

ROD RASMUSSEN: Okay. Again, if anybody comes up with any more questions on this topic or others that we covered and didn't get it in, feel free to add it at the end. We'll try and cover it then. Thank you very much, Jeff.

All right. Next slide—oops. Routing security. So we've just kicked off this work party, literally just before we got to the ICANN69 virtual conference. So we're still working on forming out what the scoping and charter of this work party look like. And we're having some interesting discussions internally on this. So this is not settled territory yet. But we want to take a look at some things that have been kicking around on our list to take a look at and various aspects of routing and their impacts, particularly on the domain name system.

Route hijacking, for those of you familiar with it, is a fairly old topic space. Problems have been with this for many years. A real brief overview. The routing system is based on various autonomous networks—ISPs and people who have IP space that interchange traffic with each other. They all agree amongst themselves how to do that. So there's a lot of trust involved, knowing who you peer with. You peer with multiple people. It's fairly easy for somebody to claim that they have some particular part of the IP address space without actually controlling it. But they can make a claim, and it can get perpetrated

throughout the routing system. I don't want to get into the details of how that works, but there are some countermeasures and things like that.

But, suffice to say, mistakes happen. Intentional hijacking happens, and things ended up getting routed to places where they weren't really supposed to be routed to. That can include, of course, the IP address base that ties to DNS servers, for example. Those kinds of combined attacks where people want to do a DNS-based attack but they need to somehow fool people into coming to their DNS infrastructure instead of the real ones may involve a routing hijack to do that. So it's an area of interest, not just from an esoteric or interesting Internet security perspective. It's also germane to people running DNS infrastructure and understanding what the risks are and what some of the mitigations are.

So those are the areas we want to try and focus on and take a look at the space on touch on that from an SSAC perspective, which we haven't really done before, and look at things like our KPI as a key tool for this and the appropriateness of that and the pluses and minuses there and, at the end of the day, hopefully come up with something that will help inform then community and may provide some information to operators who weren't fully aware of what was going on. We've seen some of the attacks that have occurred. People were educated when they happened about some of the vulnerabilities they weren't cognizant of or at least had not done a lot of work to mitigate that risk.

So that's the goal set we have there. As I said, we just started that with the ICANN69 activities. We're on pause until this is done, and then we'll get back to finishing out our goals for that work party. We'll probably have a little bit better targeting and scoping of that. Hopefully, we'll be able to get that out maybe before the next ICANN meeting occurs because we have a good long period of time between now and then.

Any questions on that? I don't see any in the Q&A. And some interesting chat but no questions. Again, as with any other topics, if you have questions, please save them and we'll try and collect them at the end.

Next slide, then, please. There we go. Just a quick update on a project we've been working on for a little over a year now. Whoops. Did you skip on there? Oh. Where are we at? There we go now. Which slide are … We're supposed to be on—yeah. There we go: the environmental scan. Thank you. There we go. [inaudible] to Internet naming and addressing. So we've been working on this environmental scan for a while. The latest update on this is that we've done a bunch of work and had some really good work done by some ICANN Fellows as well on some research to help really flesh out the various threats that we have identified to the overall DNS system. You can see the main bullets there. There are more as well.

But where we're at right now is working on particular areas with sub-teams of SSAC members doing some concentrated work on each of

these areas and also looking at mitigations and impacts on mitigations and then like.

The current working plan that we have is to start releasing documents in this space on each of these subtopics in a series of documents rather than trying to put one giant document out at once so we can concentrate on specific areas and share that with the community. We're using this internally as a tool for understanding where we have gaps in our membership or we have work where we've identified some work topic items based on areas where we haven't hand comments before or those comments might be dated and need updating. So we're using that internally, but we also want to provide that for the greater community but with full context. So the objective at this point is to concentrate on specific item areas, publish those, and then basically have a series of smaller papers that, all together, create a larger document set that will hopefully provide a baseline for looking at DNS risk areas. So that is the plan there.

Let me see if there are any questions there.

I see a question. This refers back to the last topic on RPKI. We're not sure what the main differences are yet between what OCTO did in OCTO14 and our work in the Routing Work Party. So that's actually one of our chartering questions that we're wrestling with right now because that is a good piece of work that OCTO put out. Obviously, there was a different remit there, but certainly that will be taken into consideration. So I don't know the answer to that question yet. We will find out as we go through the work to see how that differs or may add

to or augment or look at things differently than OCTO14. Thanks for the question. That is a good one.

Any questions on our progress on the environmental scan and analysis of the threats to names and addressing?

Okay. I will move on to the next slide then. Here's a list of things that we may be working on at some point here in the future. At the end of our time, if there's some other topic areas that folks would have some interest in that we may not have covered in the SSAC, we're always for idea there, too.

There were several different protocol-related things that have come up as separate topics. We've bundled those together into this meta project that we call the Evolution of DNS Resolution, looking at how things might be in the future with a bunch of different technologies and things that are emerging. So that's probably the next piece of work we take on after routing. Obviously, that's subject to change, but, as we were using that environmental-scan internal document that we have, we were able to identify and bundle those things as a project that would make sense for a work party to take on. So we're thinking we'll probably take that on early in 2021, depending on how time frees up from the other work parties we've got going.

There are some interesting operational challenges and work going on in how you manage your DSKEYs in the DNSSEC environment that have been kicking around for a while that we may well take a look at. There's a follow-on to our work on DoH/DoT and other encrypted DNS

protocol stuff. We want to potentially take a look at how that's evolving and whether there's an overload in HTTPS.

Another topic that has come up more recently is taking look at SSR data that is both collected, produced, published, etc., by ICANN Org itself but also within the overall ICANN community. Is that data that's being brought in or looked at or what-have-you appropriate for purpose? Are there some things that exist that would be useful that aren't being tapped? Are there some gaps in data and analysis that can give us a better look at how the DNS ecosystem is working and flag potential issues before they arise? So there's a bit of early conversations on this as to how to take a look at data in our ecosystem, particularly the SSR data (Security, Stability and Resiliency data). So that's still in the formative stage, but there's some conversations we've had on particular areas where our past advice is talked about. We need data for this, or the data being collected in this way is inconsistent or may not be appropriate for the use and needs to be looked at. So that's a larger topic space that we may be spending some time on.

Any questions on any of these before I move on?

All right. Next slide, please. I mentioned earlier that we have this system with the Board. It's called the Action Request Register (ARR). We actually have gone through and done a really thorough review of the items that are tracked there. Some of the things we found, which we've talked about in SO/AC leadership circles and our conversations with the Board and around the whole, "How does the ICANN world

work?", focus on what happens to things that have gone through this process and get passed along to other parts of the ICANN ecosystem. We have a pretty good idea of what happens when they get passed on from the Board to ICANN Org itself. There's a pretty thorough way of tracking that through time and how that gets implemented. But there's not a really clear indication of advice that gets moved along into the policy[making] realm or what-have-you and what then happens with that. The DNS abuse and subsequent procedures was a recent example of that, which maybe is a concrete one for you to think about. We had advice around that with respect to new TLDs, and that got passed to SubPro. SubPro looked at that and said, "Well, that's really an all-gTLD issue," and [it] got pushed back. How does that process work and how do you continue to track where that actually ends up going and eventually resolving that, one way or another? So that was an interesting bit that arose from our analysis as to how do we deal with that.

We've been having good discussions with the various parts of the ICANN Org on the implementation advice. That has been a process that's been improving over time quite a bit. Going forward, we're trying to do a better job of understanding where things are in implementation phase. I note that there's been some proposals around how to deal with advice or policy that needs to have some studies done as to feasibility and cost and things like that that we're thinking about for some of our own advice. It may need some sort process there, too. If you think about the NCAP project, that ended up getting formally costed, etc., which is capabilities we don't have within

the SSAC to do. But, for these kinds of larger-impact recommendations, it would be useful to have some understanding of resources required at least by ICANN Org, if they're involved, but potential impacts elsewhere. So we're looking at how to do a better job at providing actionable advice so that both the advisor and the advisee have a better understanding of how those may have impacts, again, looking at improving the quality of the work so that it makes it easier to actually try and get it implemented.

Any questions on that? I know it's kind of an insider thing, but hopefully those are the kinds that would be useful for other parts of the community to be able to take a look at, as we've got a fairly robust process at this point.

All right. I'm not seeing any questions or comments there. Let's move on.

Next slide. I'm going to pass this over to Julie to walk through membership stuff. Julie?

JULIE HAMMER: Thanks, Rod. And thank you for moving the slide on, Kathy. We've been doing quite a bit of consideration of the skills of our SSAC members in the last year or two, in particular in how we actually describe and gather information on the skills of our existing members, and then use that to give us some guidance on the sorts of gaps that might exist and the skills that we're looking for in potential new members of the SSAC. We've been relating to that to the sorts of work

that we've been doing within the SSAC and the sorts of work that we see potentially coming in the future.

One of the things we've done is try to make our skill survey, which all SSAC members do at the beginning of each year, much more user-friendly. We used to have an interminable number of questions, and we've now reduced that significantly and organized them into these nine categories. Within those nine categories, we only have about 47 questions. That skill survey, within the last week, has actually been posted on our public SSAC website, which we've been trying to do for some time. So, if you're interested in seeing what it looks like, that is certainly available on the main SSAC page that you can link to through the ICANN website.

Next slide, please. As I said, we've used the survey for our existing SSAC members to actually define some of the areas that we are seeking more skills on. So, while we're not exclusively seeking skills in these areas, we are particularly seeking skills in these areas.

We're also seeking to diversify the membership of the SSAC so that we have more members from different cultures, from different backgrounds, who are familiar with working in perhaps different technical environments than some of our existing members. And we're particularly interested in members who might come from Africa, Latin America, and Asia-Pacific to bring their particular perspectives, which we believe will benefit the SSAC.

Next slide, please. If there are any of you attending this call who want to have a look at the skill survey and believe that you would like to

express an interest in SSAC membership and potentially coming through our formal application process, please do read the information on our website. Contact Rod or myself or any member of the SSAC support staff by the e-mail that you see here on this slide. We'd be very happy to engage with you and give you whatever additional information you need.

Any questions, please?

Okay. Thank you, Rod. Back to you.

ROD RASMUSSEN: All right. Thank you, Julie. Just let me add that, on our membership, it's a real challenge right now since there are no in-person conferences going on around the world today. So it's a real challenge to do outreach and recruiting the kind of in-person stuff that we've done informally in the past. We do have a project we're working on with the ICANN communications team when we can get their time [slices]. They're really, really busy, as you might imagine, trying to work in a virtual world. But we're working on a project to do some outreach with them as well to various geographies, not only for those of you on the call today but for those … If you know somebody who might be an interesting person to add to the SSAC and would have interest in doing that and has the background, as you saw the skills we were looking for there, we encourage you to think about that and maybe ping that person and have them take a look and apply because it's important for us to keep fresh sets of eyes coming into the SSAC. We haven't added many members over the last year or so, and it's one of

the things that we're definitely looking to try to do. So I appreciate any networking and outreach folks can do.

Now we've reached the end of our prepared presentation remarks, etc. Now is your chance to bring in any other topics or questions that you may have—things you might want us to take a look at, questions you have about security, or issues that are going on right now. We can do our best to try and answer those on the spot or at least take them on. So I'll throw it open to folks who want to bring something up, either in the chat or the Q&A, depending on what privileges you have. I would love to hear anything else that is on folks' minds. And a softball question or two would be fine.

All right. Oh, I see a question here. "As a young person in the ICANN community, how do I build my capacity to engage actively in SSAC with time?"

There are many paths, I would say. That'd be the answer to that as far as that goes. There are the particular areas of security and infrastructure management and the like that we listed. It really comes to building up your own professional expertise in that, obviously, knowing how things work. But reading documents that we published on SSRs—not just us, obviously, but others—and attending various other conferences and things like that—IETF if you're on the engineering track of things. Or there are various security-focused conferences, where papers are presented and the latest and greatest information around threats and mitigations are presented, whether it's something that is a particular area of interest, like messaging

stuff[.] An organization like M3AAWG would be of interest if you're more of an incident responder first, where organizations from around the world that do incidence response come together. Those are really useful for building both your own background knowledge and your network. Networking is really important in this space, both professionally and for being able to bring value to any organization, whether it's SSAC or others, that you belong to because you have thee ability to reach out to others. So as much as you can there. And ask questions when you have these kinds of opportunities. If you've got an area of interest that you'd like to learn more about, this is a great opportunity. So hopefully I answered that.

Mason asked a question in the Q&A. "Has SSAC considered asking the Board to upgrade the contracts to improve the ability to mitigate DNS abuse?"

Have we considered? There may have been internal discussions around things like that, although I will say, for the most part, the SSAC does tend to avoid being too specific with its advice on how to deal with the problem, whether that's by changing contracts, updating policy, or what-have-you.

So what we try to do is explain what a particular issue is and provide some advice towards trying to mitigate that issue without dictating how that gets done. That's a bit of an art, as you might imagine, because there are multiple ways of trying to solve a problem. Sometimes it seems obvious but there may not be an obvious answer. And there is more than one way to get things done in a politically

expedient way. And we try to stay out of the politics as much as possible. So it's an area where we try and address getting things done without getting too much into the nuts and bolts of how to best do that and leave that to the Board to work out as an advisory organization to the Board.

But these things are discussed, for sure. We try to come up with ways of getting towards good results. That's the way we try and present our final recommendations.

Paul asked a question. "Does the SSAC plan to open its membership, basing it on a structure similar to the RSSAC Caucus?"

No. A quick answer. [I have] no plans on changing the membership we have right now. We've actually just finished our formal review process. Then we made some tweaks.

One thing we have done in the membership process is we've been doing a process of, as applications came in, we just considering them individually. We're trying to move towards a batching thig, where we can take a cohort of folks that come in. We're still working on implanting that, and COVID has not been helpful, as you might imagine. So we have tweaked it a little bit and then have that group considered in its entirety to see how that covers the gaps that we have because one of the things we find is that, if you do things just one at a time, you may end up not addressing gaps or have a whole bunch of people fill a gap and maybe not so much for another area. So we're trying that. So that's our latest change there.

Mason with another question: “Has SSAC considered writing a paper on how to implement trusted notifier programs?”

The DNS abuse paper we’re working on right now is going to at least touch on trusted notifiers, though I don’t think to the level that your question may imply (how to implement), but we are trying to touch at least on the concept and how that would applicable in dealing with the issues that Jeff outlined earlier on the paper. I know there are examples of those in the world today, and there are various proposals, I know, that are kicking around. They all have some merit. I personally worked with the Anti-Phishing Working Group many years on a trusted notifier program we created there. So these things do exist. They do have, depending on the situation, some good applicability. And, depending on the feedback we get from the paper we’re going to publish, we may dig into that. Particularly if there’s a movement in that direction in the ICANN space, I would certainly imagine that we’d probably be involved in that and commenting on it at some point. So hopefully that answered that question.

Any other questions or comments? Let me check the ... Thank you for answering that question more thoroughly on the RSSAC Caucus question, Julie.

Yeah, I think there was a little bit … I’m not sure where all the discussion is. We do, on occasion, have invited guests on work parties. Currently the DNS Abuse Work Party has a few to add some diversity and some expertise. I ran a work party several years ago on the Public Suffix List, and we had the actual … Well, Jothan was on that work

party. Jothan [has worked] with the Public Suffix List. He's a preeminent expert in that. It was really valuable for us to bring in a person with that kind of background. So that does happen. So there are opportunities to interact with the SSAC in an area of particular expertise. So, if you do see us working on something that you feel like you might have some valued input in, that would certainly be something to raise your hand on.

On the Routing Working Party, it's very likely that we'll do some outreach and have some invited guests. As with anything that's future-looking, do not set that in stone, but that is an area that we've already had discussion in that work party on about wanting to do some outreach and bringing in a few more routing experts. We do have a few people on the SSAC that are familiar and do those operationally, but it's always good to supplement, particularly on an area where there's as broad an impact as that.

Any other … You don't owe me a steak dinner, Jothan, but I appreciate the … I'll take a beer at some point when we can actually see each other again. We do miss seeing everybody at the meetings. This virtual thing is a challenge.

Okay. I don't have any more questions. We're getting near the top of the hour. I would like to thank all the SSAC members who were able to work this into their schedule. At one point, we had 90 people attending, so thank you very much for that. And thank you very much to our staff for preparing this and working through it and working out technical glitches on the fly as we go through. Very impressive. So

thanks, everybody, and have a terrific rest of your ICANN69 experience. Again, we are looking for more folks to apply for membership, so please keep that in mind. Thanks again. We'll see you virtually again probably in a few more months. Thanks, all. We'll end the meeting now.

[END OF TRANSCRIPTION]