ICANN76 | CF – SSAC Public Meeting
Wednesday, March 15, 2023 – 09:00 to 10:00 CUN

KATHY SCHNITT: Hello and welcome to the Security and Stability Advisory Committee public session. My name is Kathy and along with my colleague, Danielle, we will be the remote participation managers for today. Please note this session is being recorded and is governed by 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 them aloud during the time set by the chair of the session. If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor.

This session includes automated real-time transcription. Please note this transcript is not official or authoritative. To view the real-time transcription, click on the closed caption button in the Zoom toolbar. To ensure transparency of participation in ICANN's multi-stakeholder model, we ask that you sign into the Zoom sessions using your full name, for example, a first name and last name or surname. You may be removed from the session if you do not sign in using your full name. With that, I'm happy to hand the floor over to our SSAC chair, Rod Rasmussen.

ROD RASMUSSEN: Thank you, Kathy. Good morning, everybody. Welcome to the SSAC public session. I'm very glad to see you all here. It's our traditional first thing in the morning. I just like it when first thing in the morning is 9 in the morning instead of 7 in the morning. It's a little more hospitable. Most of our, I think all the SSAC members who are here in Cancun are around the table and we have a few other folks. Feel free to grab a chair at the front at this point if you want to have a space there. But we're going to dive in here with recent things that we've been doing in the SSAC and what we're doing in Cancun and looking forward to 2023.

Next slide, please. So, this is the agenda for today. We're going to start with an overview of the SSAC for those who are not familiar with it and then go into recent work we've done. We have not done any major publications lately, but we do have some correspondence documents we've done since the last meeting in KL. We're going to go into some current work parties and what we're doing in those. We are doing a little bit more on what we are working on versus what we've recently published which gives you kind of a preview which we don't always do.

And then also talk about where we are with looking at other -- I get that, thank you. With other work throughout the ICANN community. And then finish up with our membership outreach, we are in recruiting season. And for those interested, we'll have some information about what kind of folks we're looking for to join the SSAC. I will stop after and check for questions after each section or some of our members will be walking through some of the work that they're leading in our work parties.

Next slide, please. So this is an overview of the security and stability advisory committee or SSAC. Currently we have 35 members officially appointed by the ICANN board but selected through our own internal membership selection process and then presented to the board for validation and appointment. We advise the ICANN community, not just the board but the ICANN community and the board on matters regarding SSR, security, stability, and resiliency of the Internet, and particularly of the naming and addressing allocation systems. But we do talk a lot about various infrastructure elements, threats, and security issues, and a lot of different topics.

As you can see, we have lots of different areas we draw from, from our expertise. There's a whole list there. I'm not going to read them all. And that's there are even more, especially when you get into some of those broader topics, there's a lot of detail where there's many specialties. It was formed just over 20 years ago and 121 major publications and many other correspondence documents have been produced over that time period.

Next slide, please. So how this ties into the overall mission of ICANN is that we are actually part in the actual ICANN charter and we are part of the important part of that community looking at the SSR issues, which is a fundamental part of ICANN's mission. We have a formal process where we create advice, and that's through our publications. And those publications are created by work parties, as we'll call them, groups of individual experts within the SAC and sometimes invited guests. In fact, we have a couple of different work parties right now with invited guests that are not SAC members but are subject matter experts in the fields

that we're taking a look at. We do a lot of research that may be assisted by our staff or external staff.

And then we write a report based on the research and we may have findings. We usually have findings, but sometimes we'll also have recommendations. And the full SSAC reviews that document and approves it and often has feedback, et cetera, that goes through some internal basically peer review and then pushed as published. If there are any particular advice to the ICANN board, so recommendations to the board, those are formally called out within the publication. We submit those to the board, and the board acknowledges those and looks at that advice and then either and takes action on that advice.

Now, that action could be to decline that advice or to refer it to some other group like the GNSO or other parts of the community or to take and have org or some other, well, obviously org implement whatever that advice may be. There's a whole tracking system for this at this point where there's a response and request and response and understanding process. So this is all tracked through to make sure that what we publish is understood and discussed and eventually dealt with, hopefully with a full understanding of what the intent was. That's the process.

Next slide, please. So recent major publications, and these are all from last year or before, we had a briefing on routing security. We had some IDN variant work that we did and some updates to recommendations for SubPro. Those were all probably a year or more older, or within the last year, I guess, but those were all published before the last ICANN meeting in Kuala Lumpur. There's some information on this slide and

this deck should be available for taking a look at where we are and where you can find some of these things.

Next slide, please. I'm going to go over some of the recent correspondence we've had. I'm going to go through each of these docs and just make a -- this should be pretty quick, but if there's any questions, I'll check after each one of these and answer those and then we'll move on.

Next slide, please. So the first or the most recent publication, these are in reverse order of publication, I think, if I remember right, we had a -- there was a public comment that ICANN put out on a process for selecting a top level string for private use. This is actually in, this work is in response to SAC113, which was an advisory on a private use TLD. This is a TLD reserved and set aside for private use in the DNS rather than being allocated for a traditional TLD. And this has gone, there's a process going back and forth with the IATF on this. I'm not going to get into the details of that.

This particular advisory or this particular public comment was around the process that they'd outlined. And we had some concerns around the level of detail in that proposal and wanting to be able to have the community assess the actual criteria and process that was going to be determined by some experts to be determined who the experts are, so that they could come up with a string or set of strings that makes sense. And our comments were that the, that should be at least reviewed if not, or put either more detailed now or reviewed when that process is completed for what that selection criteria will be. So, there's not black

box, so to speak. So, that's what that comment was. It's pretty straightforward. Any questions on that? I'm not seeing any.

I will move on to the next one, please. We had, yes, so this is in reverse order. So, this was in December. I'm not sure if it was December, but this says 12. That doesn't mean it was in December. It was just the 12th one of 2022. We had some comments on the proposed amendments to the base GTLD RA and RAA to add RDAP to contract obligations. RDAP is the new protocol for disclosure of contact information. In the ICANN context, it's being used for domain names. It's also, I point out, being used by the RIRs for numbers as well. We had a couple of comments, and this was really in particular around a couple of things that we brought up, I believe in SAC97, around how reporting is done on requests for that data and how that's aggregated and reported by the registries.

It's a bit of a fine detail, but it's about how, if you are running multiple TLDs in your registry, how you report out the individual statistics for those registries, especially when you're running a combined system. It gets a little esoteric, but if you want to have accurate numbers, we would like to see them reported at the right level. I know that there's been, from this comment, there's been a response in the language. So that is an ongoing process.

And then there was, we also pointed out that sun setting web-based WHOIS may cause some end user confusion or difficulties if the access to contact data is presented on the web as a raw RDAP response, because RDAP is not really human consumable easily, and that would be something to pay attention to and make sure that if you're providing

web-based output for human consumption, that it be readable by humans.

And that also, if you're providing tools for aggregating or presenting that kind of data in a human-readable fashion, that you may run into rate limiting if you're doing that as a third party. So that's what we commented on there. Any questions on that? Seeing any? Let's move on to the next one. And the final one I was going to talk about this morning was our comments on the holistic review. I believe all the SOACs commented on this. I think we all had very similar response. And there were some things we did to help, we thought, was improve the clarity of the draft terms, and we also commented on the time period involved, et cetera.

These were pretty straightforward, process-oriented type things. And talking about scope and timing and how that might improve the holistic review. So as I said, many SOACs commented on this with, I think, similar concerns that was taken on board. So any questions about that? That's it for those. Let's move on to the next slide. I believe we have, name collisions analysis project. Suzanne, I'm going to hand that over to you.

SUZANNE WOOLF: Thanks Rod, and thanks everybody for being here this morning. I'm Suzanne Woolf, and Matt Thomas and I are the co-chairs of the name collision analysis project. I'm going to go through these fairly quickly. It's not a large set of slides, but also it's, everything here is in the slide deck that's part of the meeting online. So people can go into more detail, but I will not be taking a lot of time today.

The background real briefly, after name collisions became an issue in the 2012 new gTLD round over time as those results were studied and became a part of the background, the ICANN board asked SSAC to conduct studies, present data analysis, points of view, and primarily to provide advice to the board on name collisions, specifically on .home.corp.mail, and general advice regarding name collisions going forward.

A study to be conducted in a thorough and inclusive manner that includes other technical experts. A big part of what that means is that this isn't, the NCAP project is not just an SSAC work party. There are many SSAC members involved, but it's structured so that other community members are also part of the discussion group and contributing directly to the work.

Next slide, please. So far we've published case studies of collision strings, a corp, home, mail, internal, LAN, and local using DNS query data from root servers. Most of what this does is highlight changes over time, the properties of DNS queries and traffic alterations. The way we do DNS at the infrastructure level has changed in some key ways in the years since this problem was first characterized. And we're beginning to be able to quantify what has changed and what it means for our ability to detect and manage name collisions.

There's also a prospective study of DNS queries for nonexistent TLDs. The aim is to understand the distribution of DNS name collision traffic throughout the DNS hierarchy. And provide insights into where and how DNS data can be collected and assessed, which is one of the areas where things really have changed quite a lot.

Next slide, please. So far we're at the stage with the study two document where we're getting down to there's a great deal of background work that's been done. And we're getting down to the findings and eventually recommendations. And that's where you find out what all the background work really means. Some of the key findings so far, name collisions are and will continue to be a difficult problem. The case study indicates impact has increased as the internet has gotten larger and more complex. DNS service discovery protocols and suffix search lists are a continuing problem. There are so many different ways that people configure DNS and there are so many possible behaviors for DNS in the enterprise context or the internet context that all of these things are becoming more rather than less complex.

There's a notion of critical diagnostic measurements, which are defined as a way to measure name collisions by informing the assessment of the risk of delegation. These are being able to make objective quantitative observations about the behavior of the DNS. Given that we can make observations about name collision, mitigation and remediation is difficult also. Again, because the characteristics of how DNS is configured and used continues to increase. Designation of a TLD for private use as advised by SSAC and SAC113 can mitigate the risk over the long term, but not immediately. It is one of the things we've been able to verify. Existing measurement platforms could be extended to help inform applicants. There is some way to move ahead with more data.

Next slide, please. Critical diagnostic measurements, just the numeric observations we can make include things like query volume, the

diversity of sources for different queries, other characteristics that can be observed from open-source intelligence. Again, this is an area that quite a lot of background work has gone into. There's quite a lengthy discussion in the draft report. Impact or harm is determined by evaluating both volume and diversity across all of the CDMs. There's a fair amount of data, but it has to be fit together. It has to be coordinated and collated.

Next slide, please. The goal of recommendations for a workflow or a way of handling name collisions, the goal is to ensure that name collisions can be assessed. It requires name collisions to be visible if they exist. That is, if we can't observe it, we can't address it. To ensure there's an opportunity for a mitigation or remediation plan to be developed and assessed, which requires understanding the cause of names collisions such that mitigation or remediation can be developed and assessed.

Next, please. As a sample or a sort of straw man workflow that we're discussing, the diagram there has a number of stages and a number of actors. The details are still being discussed, but there's a number of measurements and a number of decision points that come into the candidate workflow here.

Next, please. There's also a role for a technical review team in the current working proposal. They would need to be independent and neutral experts. Technical expertise must include knowledge and understanding of DNS, knowledge and understanding of internet infrastructure and where other infrastructure protocols and activities intersect with the DNS, ability to review and understand the data

collected, and the ability to understand and assess risk. The proposed responsibilities are to the technical review team is there to assess the visibility of name collisions, document data, findings and recommendations, assess mitigation and remediation plans, and emergency response should that become necessary.

Next, please. There's also a role for a neutral service provider responsible for operation of the servers that will collect the CDMs. Data privacy concerns are obviously a key element here, still under discussion. There's some question about how to organize it, part of the technical review team or a separate item. If a separate team, could there be more than one? But the responsibilities of the neutral service provider as so far discussed, to operate the passive collision assessment environment, operate active collision assessment environment, where some measurements are characterized as active and some as passive, and that's where a lot of the discussion is going on around findings and recommendations. Log processing and analysis preparation to support the technical review team. And again, if there's a role for emergency response.

Next, please. So where we're at is there's been a great deal of background work done. There's a working draft of the study to report. And where we're at now is the findings and recommendations stage. And everyone who's been involved with a project of this level of complexity and scope knows that that's the really interesting part, because doing findings means assessing what is supported by the background work you've done. What findings does that imply, which in turn drives what recommendations you're prepared to make.

So we are in the process of working through findings and the implications in terms of action that might be taken, and specific answers to the board questions. We do think that we're getting close to finishing the work we were scoped to do. And the target is to have a draft report or public comment before ICANN77. And that's what we have today. Rod, I guess, defer any questions to the end. So back to you.

ROD RASMUSSEN: No, actually, if there are questions on this section, let's go ahead and take them now. Any questions about what's going on with NCAP? I know we've had a couple of sessions here already. But I know there was a discussion during the prep week, too.

SUZANNE WOOLF: So if there's no questions now, feel free to grab me. Jim Galvin has been heavily involved in this, in NCAP. And I'm sure would be happy to talk to anybody interested. Also Rod. So please feel free to ask. Thank you.

ROD RASMUSSEN: And if there is a discussion group, there is still time to join that. Go ahead, Jim.

JIM GALVIN: Since Suzanne called me out to thank me, I'll return the favor. This was not scripted, I assure you. But she took over as co-chair of NCAP late last year. And I just want to observe. I'm impressed with how quickly

she has come up to speed for someone who was not involved really in the work for the most part coming up and taking this on. And I think that's a good thing. And I really would hope that even more of our SSAC members as we come to the end would join in that effort.

And for other discussion, anybody else who's out there, certainly you can see copies of the document as it's going through the discussion group because it's posted on the public mailing list. Please do jump in and have a look in these early stages and find people to talk to. And so that we can get to closure sooner rather than later. So we would welcome input and comments from everyone. But thank you to Suzanne.

SUZANNE WOOLF: Thanks, Jim. And back to you.

ROD RASMUSSEN: Thank you. And then thank you to Matt Thomas as well. Unfortunately, he couldn't be here today. He had to go back a bit early. But he has been doing a massive amount of work, both from a research perspective and co-chairing now the work. So if you know Matt, he's another good person to talk to about this if you've got questions. Let's go ahead then and move on to the next topic. That is updates on ongoing work parties with some more detail than we typically do. We're going to have the next few slides are going to cover the second, third and fourth bullet points here, which is because we already talked about the first one in our previous section there. We have three work parties going on and we'll get talk to all three of those.

And we have ongoing work there and put a plug in for today's DNSSEC and security workshops. We have, I believe, two or three sessions of those later today. They're open and on the agenda. I encourage you to attend. All kinds of interesting stuff around what's going on with DNSSEC and other security issues. And then we will talk about membership outreach towards the end. We do have a standing membership committee that interviews candidates and does that kind of work. Next slide, please. DS Automation. Peter, I believe I'm handing this over to you.

PETER THOMASON: Yes. I hope you can hear me well. I'm Peter Thomason. I co-chair this work party together with Steve Crocker. And so the reason why this work party exists is essentially shown on this diagram. It perhaps looks more complicated than it really is. I'll explain it quickly. So let's say you have a domain name and you booked a package with your DNS provider and includes DNSSEC signing. Then the DNS provider signs your stuff for you. So for each IP address that you put in your DNS zone, there will be a signature and all of that. But it's no use if it cannot be validated. And to validate it, you need to convey some of the DNSSEC key parameters to the parent registry so that your domain becomes part of the global chain of trust.

That's called DS initialization because you have to put a so-called DS record. Another way that usually works involves a manual step by the domain owner. So if you look at this diagram, the registrant will usually have a contract with the DNS operator and then the DNS operator as a first step, that's the arrow with number one, will put the signatures on

their server. Then the registrant has to interact with the DNS operator and fetch these cryptographic parameters to then transfer them upwards to the parent or TLD. And then how exactly that works depends on the specific circumstance. So it can be directly or through a reseller. And there are different kinds of web interfaces and different formats that can be presented.

So you can't even as a DNS operator give a very general advice to the registrant how to do it. It's kind of a process that depends on the case, how exactly it works. Then once the parent registrar does have the information, it's moved on to the registry, which has the TLD database essentially. And then they put that DS record that you gave them into their TLD server. And that establishes the connection in the global chain of trust with the zone that's hosted at the DNS operator at the bottom. And the complication here is that the two orange arrows are often manual and it's not always the same kind of steps that need to be taken there. It's a highly specific process. And that's where it seems that there is an obstacle to DNSSEC adoption.

Next slide. So actually, the SSAC has thought about this stuff a bit. So we looked into various studies that are available and evaluated how DS management is done at various registries and registrars. And there is, for example, a study that shows that even when a domain owner whose zone is signed tries to go through this process of getting a DS record deployed, they don't succeed in all cases. So actually, in 40% of cases, it's not completed at least, that is what the SIGCOMM paper from 2017 showed. So that means even when it is attempted, it doesn't always work and that kind of screams for automation.

And second, there is a bunch of things in the infrastructure that are automated from a domain owner's or website owner's perspective. Not only the DNS stuff is automated to a large part, you don't have to take care manually of replication for example, but the operations of the web server is also automated often. You don't have to take care of getting a TLS certificate. That's usually something your web host does and it gets renewed without your interaction. With DNSSEC, somehow it's different. You have to do something. And so it appears that the way that DNSSEC is operated today is kind of in contradiction to what a registrant could expect in terms of what the service is that they get provided.

Now, interestingly, there is a bunch of protocols that were developed in the IETF. And the footnote here is the references for the RFCs. And it's actually possible to automate this step that we just discussed in the previous slide. And I'll spare you the details, but you can ask if you're interested. And so essentially it works by the DNS operator on the child side publishing some information and the registry can proactively discover that and digest it and then insert this DS record on their own. And there is a few problems of how exactly to do that, but there are such mechanisms. And it turns out some ccTLDs and some registrars and some DNS providers have implemented that, but so far it's not very broadly adopted. So it's like 10-ish TLDs perhaps.

Next slide. So when you turn this stuff on, actually, there is a bunch of technical considerations you have to settle before you have a process that works well everywhere. So for example, when you are a TLD, let's say you are.se and you are implementing this automation protocol, then you have to think about things like what happens when the

domain owner has put a lock on the domain. The lock often is there to prevent accidental hijacking or change of the registrant name or accidental or malicious change of the name server records in the delegation. So now when the domain is locked, is it fine or is it not fine to automatically manage DNS using those automation protocols that I just talked about? So that's the question you have to decide.

And actually .se, for example, I think has decided it's not fine. So they disable the automation when there is a lock and there's other registries, I think Czech Republic, for example, who have the opposite opinion. And it seems important that the registry community and the ICANN community kind of establishes consensus on such issues. There are some other issues and they are not showstoppers, but to live up to the registrant's expectations that things work roughly the same regardless of what your TLD is. It seems like it would make sense to discuss these questions more broadly and come to a recommendation overall.

So SSAC is trying to take these things into consideration and consult with various concerned people and so on and so forth. Also, there is a few questions that we have asked ICANN org, specifically whether there is any things to take into account in the registry contracts, for example. And so we're considering the response that we received last week also in the document that we're preparing. So there will be an advisory document probably this year and it will shed some light on this whole problem space and most likely come up with recommendations of what to do, what not to do and what perhaps to find more consensus on at some later time.

Next slide. So I'm almost done. So this is actually kind of the final slide. So remember the first slide, which had the complication on the left-hand side where the registrant had to do stuff and then interact with the DNS operator and the registrar. Now you can see that that layer is gone. And instead in the automation model, which is how it could be done when this gets adopted more broadly, is that after signing the zone, which is step one at the bottom, the registry can proactively talk to the authoritative DNS server of the child, do some kind of exchange, and then put the DS record directly into the TLD server. That's step three.

And after that, the registrar has to be notified so that they know that the delegation has been updated. That's step four. And very similarly, if the registry doesn't do the scanning, next slide, then instead of the registry doing it, it's also possible that the registrar does it. So then the big arrow will be on the left.

Next slide. there it is. So that's an illustration of how it could work. And I'm happy to take any questions. You can also talk to me later or shoot a second email. Thank you for everybody who helps working on this. Particularly, I want to take the chance to thank the SSAC support staff and Steve Chang, who works on this a lot too.

ROD RASMUSSEN: Do we have any questions? I'm not seeing any in Zoom. Thank you very much, Peter. And looking forward to this getting completed. It's an important challenge in the deployment of DNSSEC at the second level. And hopefully we can help alleviate some of that challenge. Next slide. I'm going to hand this, I believe, to Barry.

BARRY LEIBA: Yes, hi, I'm Barry Leiba, I co-chair this Evo of Reso work party, along with Russ Housley, who's sitting across from me. Recently, particularly, so Rod talked about how a lot of our documents make specific recommendations to the ICANN board and also give information to the community. We've got a number of documents recently, particularly that just are focused on giving advice and information to the community. So SAC121 for routing security, we did. SAC105 for Internet of Things. And SAC109 talked about some DNS resolution protocols that add security and privacy to the system by doing it over encrypted channels. That document, 109 for DNS over HTTPS and DNS over TLS, talked about issues with the privacy and encryption from the points of view of different stakeholders with different perspectives.

We got together to discuss whether there was anything more we wanted to say about it and wound up converging on looking at different DNS resolution mechanisms, or I should say different name resolution mechanisms that don't necessarily go through the DNS, such as block chain resolution mechanisms that store the name resolution information in block chains.

Some of the background on this is that, as we see it, domain names are being increasingly less relevant to end users directly. You are now in the mode where you type something into the combined search and address bar in a browser, and the browser figures out what to do with that string you typed, whether to resolve it through DNS or to do something else with it, whether to do a search on it. You're getting a lot more of your results from searches.

You're opening dedicated apps that do what you want them to do without the domain name being in your face. You're scanning codes, QR codes, on the sides of buses and billboards or on the tops of restaurant tables or on museum exhibits, and you just rely on the result of scanning that QR code to get you the relevant advertisement or the restaurant menu or the description of the artwork that you're looking at, and you don't really know how it got there or what domain name was involved. But the domain names are still critical at a lower level. The protocols underneath still use the domain name.

So what are the implications of all of this to alternative mechanisms for DNS or for name resolution? It's hard for us in SSAC to stop saying DNS, even when we're not talking about it. So what we intend to do here is come out with a report that gives our view of what the implications of this are to the user experience, which is actually a significant part of the security and stability of things. Users need to get what they expect when they do something, and not getting what you expect can be a security issue, can be a privacy issue, can be a stability issue. So the document that we're writing may or may not have any formal recommendations. It's aimed at information. It's aimed at suggestions to software developers about what to consider. And that's where we are. We intend this document to be not at a super technical level. It's going to be something that the community can really digest as a whole.

We've gotten a lot of good feedback for the documents like that that we've been doing recently. And people have said that they tend to be useful. So we hope we come up with something useful here. And we're hoping to get this done this year. So that's really all I have to say. Russ, do you have anything you'd like to add?

RUSS HOUSLEY: No, you did great.

BARRY LEIBA: All right. Any questions? Any comments? Okay, thanks. Back to you, Rod.

ROD RASMUSSEN: Okay, thank you very much. I don't see any questions in the Zoom room. Please don't be shy, folks. This is a chance to ask questions as we're actually doing some of this work a little bit more in depth, kind of behind the scenes of what we're working on. The next work party is being run by Casey Claffey. Casey was unable to attend, and I know she has a conflict this morning. So I'm going to go ahead and cover this for her. But she is leading this work party and has got myself to a great start. This is a particular issue around the management of name servers. That's what NS stands for, hence title. And it's how they're managed in a particular circumstance. I'll just run through the problem area and what has occurred over time.

And this is actually a subset of other problems with when you run into a lame delegation or an expired domain that hosts name servers in particular. And we're actually doing a work party work yesterday and research found that we've talked about this problem several times in kind of a more generic sense. And what we're looking at here is what's kind of evolved over time to deal with this particular issue, which is creating its own set of issues. So we're looking to try and see if we can

bring some thoughts to being more consistent and doing a better job of solving this particular problem.

So if you have a domain name and you're using name servers that are controlled by a third party and that domain expires or something else happens to it that takes it from being able to be resolved on the internet and providing name service, that creates an issue where if you're looking up the domain name that's using those name servers or the missing name server now that you won't get an answer. If you have other name servers that are defined elsewhere on other domains, it'll still work potentially.

But there is a problem there. And there are policies in place to prevent the removal of those domains when there are name servers on it where other domains depend on it. But registries can only see what's in their own zone. They can't see if it's across TLDs. There's an issue there. So a workaround that's been that has grown up over time and practices grown up in time is that the registrars will create a new what we call a host object, which basically is a name server entry.

And they will use a non-existent domain name in another TLD operated by a different registry, which allows them to remove that name server entry from their own TLD because the domain name itself is no longer in their TLD. So the issue with that, of course, is that someone could register that domain name in that other TLD and do something with it. Set up a response where, so for example, if you had a bank that you had one of its name servers was on a name server entry, the domain had expired, there was no different TLD, somebody could go and register that domain, there's a different TLD, set up a name server to respond to

answers for your bank's domain name and send you to a fake bank instead of the real bank. And that's a problem.

And this has been a known issue, but there was a recent study done that showed that over half a million domain names have been exposed to this as at least that many, it's probably far more because that was on a fairly limited data set that it was looked at. And I believe that number was well over 100,000. I don't remember how many over 100,000 words, some activity like that had happened where a domain had been one of these zombie domains, if you will, had been registered and name servers had been set up on it to do interesting things.

Next slide, please. So this was presented at a recent, I believe, DNS symposium, and ICANN actually put out a advisory on this to the registrars. And that was recently in January. So that was, and you can see in the bottom, that is the name of the paper, the risky business paper. Next slide, please. So what we're doing in this work party is excuse me, working on the scope of this and definitions for the risk. So it's well understood and then approaches to actually deal with those issues and get rid of this hijacking, this particular hijacking risk, especially when there's different approaches being taken by different registrars and how to do this and some are more problematic than others.

And then being able to publish that along with hopefully some recommendations, and we'll see where that leads to discussions around updating operations, et cetera. There's some references there to the paper or papers that looked at these issues. There are various SAC documents that talk about this problem in the generic sense, all the

way back to, I believe, SAC four or five. It's been around this long, but this is one of those byproducts of the interdependency of the DNS upon other parts of the DNS, where you have to be careful to make sure that you understand those interdependencies. You can almost think of it as a supply chain exposure problem, if you're familiar with those kinds of issues.

That is that work party. We hope to get that one, because of its scope is fairly narrow. We do have several outside experts we brought in on that, mainly to do the operations at registrars. So that we have people who are actually very familiar with this, and we can have discussions about how to, what some best practices are and solutions to help solve this. So we're hoping that this will be done within the calendar year. Any questions on this one? Seeing none.

Next slide, please. So this gets back. Now we're going to talk about what else we're doing, just in general. There's the top three priorities. I've heard it's really one plus three others, but these are the ones we're looking at from an SSR perspective, where there's involvement, DNS abuse, access to registration data for gTLDs in particular, but also SubPro or adding new gTLDs in subsequent rounds going forward, which everybody's talking about. Now, these are the issues that are the top of mind for pretty much all of the ICANN community. There's some references here towards recent reports and documents that the SAC has published on these topics. I'm not going to read these all off. It's quite a few areas.

We've been participating quite heavily ever since the EPDP was started, for example, and with the rest of the community on how to address the

access to main registration data. And we'll see what happens with the latest plans to implement an access system. We have obviously been working very hard with the NCAP project on one of the key areas for being able to launch a new round or access to new gTLDs. So that's been ongoing work. And we've recently commented on other aspects of that program, including concerns around areas like DNS abuse. And that's an area that we're looking at as a DNS abuse itself.

Obviously, there's a lot of work going on. We're going to be hopefully hearing soon about the outcome of negotiations between ICANN and the contracted parties there. We had SAC115 was talking about some of the areas that I know they're working into this new agreement. And there are things probably to build upon there from a policy perspective.

We're talking about those kinds of issues amongst the SSAC leadership and how we might be able to coordinate efforts there moving forward once we see what the new baseline is with DNS abuse out of those contract negotiations. So those are all things we're tracking. We obviously with the NCAP project, we're still doing a lot of work on that third item there. And we are definitely tracking along on those other two and having discussions with other groups about what next steps and things may be and talking with the board around priorities and also the strategic, how that fits into ICANN's strategic vision.

So those are the things we're doing along with the community. Any questions on that before I move on? Oh, there are two things in the Q&A pod. Oh my, RPKI and quantum computing. That has nothing to do with what we're talking about today, but we did have a publication about quantum. Oh, I'm sorry. The other Q&A pod, there's a question.

As with RPKI, if a quantum computer were to become powerful enough to break these cryptographic algorithms, the digital signatures used in DNSSEC could be compromised. This could potentially allow attackers to spoof or manipulate DNS responses, redirecting users to fraudulent websites or intercepting sensitive information.

UNKNOWN SPEAKER: Slow down.

ROD RASMUSSEN: I don't believe we have translators.

UNKNOWN SPEAKER: No, but it's still hard to understand.

ROD RASMUSSEN: Anyways quantum crypto is an issue. And in fact, we are on top of it. The National Institute for Standards and Technology is developing quantum resistant cryptographic algorithms. And we and some other organizations have been involved in that. And at least had comments explaining to them what we would need for an algorithm to continue to work with DNSSEC, because DNSSEC has some size limitations that don't apply to other contexts like SSL certificates. So they are coming up with new cryptographic algorithms. They're in the process of standardizing them probably later this year or earlier next year.

Down the virtual hall in the IETF, we're looking at these things and we will probably add new cryptographic primitives to DNSSEC that are

quantum resistant so that people can start implementing them. But this is a very slow process. The entire rollout of defining these and implementing these and rolling the keys is going to take probably the better part of a decade, which I think is consistent with still being safe because nobody really expects quantum computers to work by then.

UNKNOWN SPEAKER: I'd like to add one thing to that.

ROD RASMUSSEN: Okay, we're actually running out of time.

UNKNOWN SPEAKER: Just posting a link. So there is a statement from the ICANN office of the CTO in the matter of post quantum algorithms and DNSSEC. I encourage you to read it and I'll post it in the chat.

ROD RASMUSSEN: And there was a SSAC document about this as well. The Q&A, we'll take about care of that. There's another question in the Q&A pod around DDoS against RDAP servers. We'll answer that offline because we're getting towards the top of the hour. And I want to make sure we cover the rest of this good, interesting questions, but not quite on topic for today. So the next slide, please. Thank you. So some ones we have in the hopper for potential work parties coming up are long term implications of namespace expansion. What if taking a look at theoretical limits and where things may end up, if you keep expanding

and adding TLDs at the top level, what the ramifications may be looking at data that's available for from ICANN and other places.

What are some of the gaps? What are some of the things that we can be using more effectively to understand the DNS system and related infrastructure? And then one on the looking at the technical, political, technical implementations of a removal of a TLD, force removal and what that would what kind of impacts that could have. So those are the things that we have. There are a few other things that are not quite as formed that we're looking at potentially around DNSSEC and some other things. Next slide, please. I'm going to hand this over at this point to Julie Hammer. Julie, please take it away.

JULIE HAMMER: Thanks, Rod. And, good morning, everyone. Julie Hammer, Vice Chair of SSAC and in that capacity, I'm also Chair of the Membership Committee. And in that role, we rely greatly on our SSAC member skills survey, which asks of all of our SSAC members details about their skills in nine various capacities that you can see on this slide. And we have a very detailed skills survey comprising 41 questions in these nine capacities. And we use that together and aggregate the skills of all of our SSAC members and review them manually to see where our skills gaps are.

Next slide, please. And those skills gaps, once we've analyzed them, inform some of the new skills that we wish to reach out to the community to meet in potential new SSAC members. So those skills that we would like to seek in potential new members are shown on this slide. Obviously, they don't all come in a single person. But, of course,

in addition to this, we're always interested in members who might supplement our existing skills that are mentioned in the more detailed SSAC skills survey.

We also look at a range of diversity elements to make sure that we are in a position to represent as many perspectives as we can in our consideration of technical issues. And in that regard, we're interested in increasing our representation with members from Africa, Latin America and Asia Pacific. So if you are a technical person with a security expertise from one of those regions, do please think about reaching out to SSAC regarding potential membership. We're also looking out for members with an academic background. So those of you in academia, please think about it too.

Next slide, please. We do have a particular sequence in which we reach out to the community and solicit new member applications. And it's shown on this slide here. We're currently in our outreach and application period for another month or so through to the end of April. We particularly like to receive new member applications and then the membership committee looks at them as a group of applications and engages with potential candidates and then puts through recommendations to the SSAC itself and subsequently to the board for appointment of appointment of new members. So that does not mean that we would not accept an application at any time of the year. And if it were someone who was meeting one of our important skills gaps, then that may well be dealt with out of sequence. But this is a general annual sequence of events for membership applications.

Next slide, please. So if you think you have got skills that would be of interest and used to the SSAC and you are interested in potentially being a member of SSAC, please have a look at our public website where there is a great deal of information about how the SSAC functions in our operational procedures, but also details about the reports we've published and the types of focus that we work on. Contact any member of SSAC support staff by sending them an email. But also while you're there in Cancun, you can look around the table and see all of the friendly faces of the SSAC members that are there in person and I'm sure every one of them would be more than happy to have a chat in the corridor and give you a little bit more insight. So back to you Rod.

ROD RASMUSSEN: Thank you Julie. One next slide please. One last point because we're at the top of the hour. Julie and I have already commenced our sixth and final year as chair and vice chair. We will be having a process to elect a new chair and vice chair in May. So news to come on that at some point here shortly, well next couple months and stay tuned for that. It's been a pleasure doing all this work and we have till the end of the year, the calendar year before we hand the reins over to someone else. That's it. Any final questions, comments, et cetera? Don't see any there. Thank you very much. Appreciate all your, we had a pretty good attendance online as well as here in the room. Again, thanks for being here and enjoy the rest of your time here in Cancun. We are adjourned.

[END OF TRANSCRIPTION]