ICANN79 | CF – SSAC Public Meeting
Wednesday, March 6, 2024 – 10:30 to 12:00 SJU

KATHY SCHNITT: Hello and welcome to the SSAC public session. My name is Kathy and along with my colleague Danielle, we are the remote participation managers for this session. Please note that this session is being recorded and is governed by the ICANN expected standards of behavior. During this session, questions or comments submitted in chat will only be read aloud at a time permitted by the chair. The session includes automated real-time transcription. Please note the transcript is not official or authoritative. To view the real-time transcript, 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 Zoom sessions using your full name, for example, first name and last name or surname. You may be removed from the session if you do not sign in using your full name. And with that, I'm happy to hand the floor over to SSAC chair, Ram Mohan.

RAM MOHAN: Thank you so much, Kathy. Welcome. I remember a long time ago where the SSAC public sessions would be scheduled at 8:00 AM on Thursday mornings. So typically, there were like only a few SSAC members and almost nobody from the public. So I'm glad that we've managed to change the timing and we're also looking to see if we cannot standardize on doing our public sessions Wednesday mornings, you know, schedule permitting. I would also invite, before we get started, I

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

would invite members of the public who are here to come join us here around this big U. There's no need for you to just be sitting in the back. You know, we're going to run this in an informal way. There are microphones here. There's a microphone there as well. But if you'd like to come to the front, please do.

So with that said, let me start off with just what we intend to cover on the next slide. So we have a fairly full agenda for this 90-minute session. A very brief overview of the SSAC for those who do not know who we are or why we are. And then over to Tara for new member outreach. And then to Suzanne and Matt to provide our view on the named collision analysis project, NCAP. Then to Jeff Bedser for DNS abuse. To Barry Lieba for a briefing on SSAC 123. And then I'll provide an update on the current work that the SSAC is doing. And then we'll be open for Q and A from anybody who may have such questions. So with that said, let's go to the next slide.

Just a quick overview of the SSAC. Before we get to this slide, just on the previous slide, I wanted to start with just a little bit of, you know, what got the SSAC started. In the 1990s, there had been concerns about attacks on infrastructure of the internet from various players on the internet, and there was a concern about that that was generalized. And then in 2001, 9/11 happened, and that was a different kind of attack, but it prompted that question that had been discussed in different circles. It brought that question to the fore. And that question was, if there was a concerted terrorist or other attack on critical infrastructure of the internet, the domain name system, how would we even know that there was something like that? How would we know what to do? How would we know what best practices are? And that led at the ICANN meeting in

2001 and Marina Del Rey to have a group of us to be brought together to think about, should we form a group of experts who have knowledge, background, operational, and other expertise, who can look at emerging trends that deal with security and stability, and provide advice to the ICANN board and to the community on matters that relate to security and stability.

So if we go to the next slide, in 2002, the SSAC formally got established, and it was in line with ICANN's own mission, and if you look at our own charter, it's the fundamental thing that we do is to advise the ICANN community and the ICANN board on things that primarily pertain to security and integrity of the naming and address allocation systems. So you will rarely see the SSAC speak about policy matters. You will rarely see us opine on those kinds of things. We try pretty hard to stick to our knitting. We try pretty hard to provide advice on things that we know about, and we are perhaps even reflexively very clear in saying we don't know on things that we don't actually have enough of a clue about. So that's kind of the way it is structured. Next slide, please.

What we actually do is we engage with the internet technical community and key DNS infrastructure operators, and together we come together and we articulate what security requirements are evolving. We're generally forward-looking. When we analyze risks and threats to the naming and address allocation services, we're usually not the group that is going to be responding to the latest DDoS attack or the latest malware thing that is going around on the internet. That's typically not what we look at. We try to look at more systemic, the kinds of issues that are more pervasive and that the community and ICANN ought to be specifically informed about. And we coordinate. We have

lots of friends in the technical community, IETF, the RSSAC, RIRs, other name registries. So we coordinate, collaborate with all of them, as well as in the security, on the security side of things, folks who are in, you know, anti-abuse in areas like that. We work quite closely with them as well. With all of that done, we form work parties and we arrive at recommendations and conclusions, and we then inform the ICANN board. And we also provide advice to the community as a whole. Next slide, please.

This gives you a, without actually giving you every name, but this I think is actually all of our members of the RSSAC were spread in some parts, were concentrated in some parts around the world. Tara will speak a little bit more to that. But we come from diverse backgrounds. But in general, the unifying theme for all of us is a very deep and keen interest, expertise, and ability to contribute on matters that have to do with security, stability, and resiliency. Next slide, please.

This year, there's a new leadership team for the RSSAC. You see all of us here. The team is arrayed around me here. But we also look at the leadership team as a team that consists of not only the RSSAC members, but also our incredibly qualified and just fabulous staff members. So together, we really form the administrative part of getting the SSAC to organize, getting the SSAC to do its work. And you will see below all these photos, you will see, you know, from just the team that is arrayed there, some of the skills that we bring. So it's usually qualification into the group tends to be based on the skills you have and the skills you bring, rather than the positions you hold. So that is a quick overview of the SSAC, and I'll pass it on to you, Tara.

TARA WHALEN: Thank you very much, Ram. Can I have the next slide, please? Oh, would you like to? I can talk to our publications unless you would like to. Please, go ahead. As we are going to be talking, I guess, about the breadth of expertise. Also, yes, Tara Whalen, Vice Chair, by the way, if I didn't get my introduction in. This gives some idea for people who want to see the types of work that we engage in. It gives you an idea of the types of topics that we have here. Here we've clustered them into some general categories around security and DNS management, registration data, and our catch-all of other. So this will give you a place if you want to get oriented to the work we do, or just generally some things you might find of interest if you are excited about technical topics in the ICANN space. You could start here and dive in. And also shows you, again, this sort of breadth of skills that are needed to tackle the large amount of guest threats and issues that come up in this space require this amount of advisory. Next slide, please.

Oh, wait on, wait. Go ahead, Barry.

BARRY LEIBA: I just want to point out that looking at the number of the publication numbers on all of these, this slide is kind of highlighting a lot of older publications, but we've put out about 125 publications altogether. And so routing security is a relatively new one, as are a couple of the others. We're constantly doing more work and putting out more documents and having a great time with it.

TARA WHALEN: We are indeed. Thank you, Barry. And with that, now I guess we've come to the new member outreach part. So I'm very happy to talk to people as the membership committee chair. I want to be able to connect with people who might actually want to join SSAC now or in the future. Of course, the internet itself is evolving and evolving fast, and the SSAC itself also evolves and changes. So to that end, we are looking to diversify and sustain and replenish our membership. So today I wanted to talk about some of our diversity goals and what it is that we're looking for in terms of skills and expertise, why you might want to join us, and how to do that. Next slide, please.

So I'll note first for geographic, Ram hinted at this earlier. We are scattered around the globe, but we do have a bias towards Europe and North America. So given that we are attempting to talk about global internet issues, then we feel it would be better to have more representation from around the world. And we are definitely short on qualified members from Africa, from Latin America, and from Asia Pacific. And as well as geographic diversity, we also are looking for people with skill sets from the academic and research backgrounds. We are often tackling very large-scale issues that often involve measurements, data sets, and analysis. And for us to get these very deep analysis for our advice, we need very strong expertise in these areas so that we can sort of get the right information out of the data and present it accordingly. And when I mention that, I also want to say we sometimes get people at this point an early stage of career. So if we have people who are perhaps graduate students, new graduates, at the earlier stage also have lots to contribute. And we've had some new members in this group. So you need not feel that you have to be at a

late-stage career person in order to have expertise to provide to us. We very much encourage people at earlier career stages so that we're all sort of not aging out, for example, at the same time. And also replenishment, people do come in with new approaches and new skills and new tools, which are absolutely valuable for us to do our work more effectively and more efficiently. So we're very much looking for that as well. I'm also going to note that we are a lot more men than women at this point. So I'd be very happy if we could address the gender balance for sure. Can I have the next slide, please?

So we have a word cloud here of some of the skills. You don't need to read that literally. It's really just meant to be a representation of the breadth of areas. We had some indication of that in the list of the publications. But it shows really the breadth of different topics. Now we are a technical group at core, but of course the technologies themselves cover a wide variety of areas, lots of different topics. There are operational matters. There are theoretical matters. Technology is changing all the time with new threats. So if you have any kind of skills that you think would be relevant for the topics we cover, then we can make use for you within SSAC. There are broad skills as well that can be presented as well as technical topics. So if you, as I mentioned, data analysis for example or risk management are things that are very valuable, as long as they apply to security and stability issues in the ICANN context, then that will be a very valuable contribution for SSAC. So you can have a look at some of our advisories and think about how you might have contributed to one of those documents. There's an area perhaps of your expertise where you said that is exactly the sort of advice I would have been giving. Also you might notice a gap. There's

something this document was missing that I could have put in there if I had been able to contribute and that's exactly the sort of thing that we're looking for. Can I have the next slide please? So we hope that perhaps also because our meetings have been open and you may have looked at some of our projects, some of our engagements, you are excited about some of this work. You may already be coming here with excitement about engaging with ICANN. If you like to engage on security matters, interesting, thorny problems of security and stability on the internet, if you'd like to interact with other experts in the field on these topics, if you're interested for your career-wise in deepening and extending your expertise and having international visibility in some of the security forums where a lot of these topics are discussed, if you want to advise the board and advise the community and sort of help build a more secure and stable internet, well please apply to be a member. We're absolutely welcome you to do so and help us in our mission. We now have an online application process which is on the SSAC, on the website within the ICANN space. So that is www.icann.org/n/ssac to apply. And you can talk to the SSAC staff for questions or please talk to us. We're very friendly people and we would like to talk to you. Next slide, please.

So one of the ways in which we're trying to contact or connect with people is we're doing an experiment this time as a new SSAC open forum. So we have an informal drop-in session where you can talk to us directly. We'd like to talk to you in depth, address all of your questions about our current projects, work that we might be engaging in in the future. If you want to ask deeper questions about becoming a member, find out more about our mission. If you saw something this week that

you want a little more information on, then please come and talk to us on Thursday in Block 1 in Room 103B. We would love to see you there and to talk to you. Thank you.

RAM MOHAN: Thank you, Tara. And for those of you who would like to informally come and chat, if you cannot make the Thursday session, some of us will be hanging around after this session is over and come and, you know, meet us and chat with us. So thank you for that, Tara. Let me move on to the Name Collision Analysis Project over to you, Suzanne.

SUZANNE WOOLF: Thank you, Ram, and welcome, everybody. I'm here with my co-chair, Matt Thomas, to discuss the Name Collision Analysis Project Study 2 Report. We've presented on this material several times this week, and we have plenty of people in the room that are familiar with it. So we're going to kind of keep it moving. But as Ram noted, the draft report was published a few weeks ago, public comment period ended earlier this week. But I do want to make sure that we thank the discussion group, the cross-community discussion group that produced the report. A lot of work, a lot of effort to reach consensus on some difficult issues. So thanks to all of them. And we also note that we tried to work by consensus as much as we could. So the contents here are the group's best assessment of the issues and can include the best guidance we felt we could offer. Next.

So as a quick recap of all the reasons for this work, how did we get involved in this? First of all, what are name collisions? And one simple

metaphor to explain it is if two houses share the same address, it becomes really hard to know where mail or delivery should go. And things being delivered to the wrong place, there can be a lot of very serious bad consequences just to getting the right thing to the wrong place. And in the DNS, this can happen in the sense that name collisions occur when a domain used in the global DNS namespace is also used in a different namespace. You know, often a private enterprise, you know, a corporate IT department or something along those lines. But it can happen in a lot of places. It can happen in a lot of ways. And generally what happens is users or software or other functions may misinterpret data they get and have trouble doing what they're supposed to do. Next.

And why is it relevant for the new gTLD effort? Because this is part of the general planning for new gTLDs. There have always been name collisions with us and there has always been some risk of unintended consequences. You know, businesses have used labels for internal TLDs in private namespaces that could leak to the global internet. There's always been that, but introducing new gTLDs increases the probability of name collisions just because there are a larger number of potential names to collide with. You know, and there's more risk that there's unintentional overlap with names already used in private networks or internal naming systems.

And the other thing that we have to make sure we take into account in the new gTLD round, measuring name collisions is difficult, largely due to evolution of technology and network infrastructure. There were certain things that went into name collision analysis assessment in previous new gTLD round, but it's the internet. The technology has

changed and legal and regulatory constraints have changed. So, for instance, privacy enhancements in the DNS and alternative naming systems have made the DNS landscape more complex and measurements more difficult. So what follows is how the discussion group tried to frame those issues and some recommendations on how to manage them. And Matt's going to go on into what exactly we had to do and what we came up with.

MATT THOMAS: Thanks, Suzanne. Next slide, please. So ICANN board asked the SSAC to convene into the NCAP discussion group, which is a community-oriented group, to provide some advice around how a secure, stable, and predictable framework can be constructed to assess name collisions going forward. This was notable because name collisions are clearly a factor in the next round of gTLDs. So to that end, the board asked SSAC to provide some data opinions, point of view, findings, and recommendations around three main topics, the first of which is a specific advice around Corp, Home and Mail which, as you may know, are holdovers from the 2012 round, as well as answering a series of board questions. I think there were nine board questions in total that provided a kind of overall generic advice around name collisions and how to look at them going forward. And as I mentioned earlier, all of that research and discussion was done with the community in the NCAP discussion group in an inclusive manner with participants from a wide variety of backgrounds and interests all working together, hopefully to provide our best efforts here to provide the Study 2 report out. Next slide, please.

So the three studies that are contained in the Study 2 report provide some interesting findings for the recommendations in the Study 2 report. But the three studies here that I'm going to highlight are the first ones looking at collision string specifically on Corp, Home and Mail. But we also looked at LAN, local and internal using DNS telemetry data from A and J root servers going back to 2015. That study was insightful because it provided a longitudinal view of name collisions over several years. And one of the clear measurements from that was that name collisions have increased during that time and are continuing to persist. Out of that work also was the identification of what the group has called critical diagnostic measurements and this is a set of quantitative measurements that can be applied to DNS telemetry data that are useful for potentially measuring the impact of name collisions.

Now, one of the second studies was looking at a perspective study of DNS queries for non-existent top level domains and this study was really focused on where and how DNS data can be used for risk assessment purposes for name collisions. As Suzanne mentioned earlier, the DNS has changed significantly since 2012. We have things like QNAME minimization that are impairing the observation of the names being sent and collected at the root servers for name collisions. We have other technologies like aggressive NSEC or local routes and now even the introduction of large public recursive resolvers that we really didn't have in 2012.

So this study took a look at various different root servers and was able to provide some insights in terms of where and how name collisions might be assessed using that data but also was able to look at data from recursive resolvers to compare what name collisions look like there. So

the overall findings from this study really helped provide some guardrails around where and how data can be collected and what the pitfalls and benefits of certain observation spaces might be.

The last report was looking at the root cause analysis of new gTLD collisions that were reported to ICANN since 2012. There is a form that has been available on the ICANN website for those experiencing name collision issues to report it to ICANN. And this study was taking a look at the 47 different reports that ICANN has received since 2012 to understand what was the root cause of those name collisions being reported, what kind of impact or severity were those name collisions, and how could that information be used to provide guidance going forward. All of this together really formed the core aspect of what the discussion group has highlighted in the Study 2 report and that is name collisions are a risk management problem and that name collisions are here and they still persist a threat to the DNS security and stability. Next slide, please.

So what's coming out of the NCAP Study 2 recommendations really focus on a framework for assessing name collisions going forward. As I mentioned before, this framework is tailored around being a risk assessment process. And that risk assessment process needs to be conducted by a organization that is being prescribed or called the technical review team. And this is a group of individuals who are there to evaluate the new gTLDs based off of the quantitative as well as the qualitative data that they're collecting on those collision strings.

Now to get that data, one of the things that the report highlights is since the DNS ecosystem has changed significantly since 2012, the way which

we get that data also needs to be altered. And we will talk about that a little bit more going forward. But essentially in order to collect better data since the data that was used in 2012 is either A, not available or is highly impaired, things such as delegating the TLD into the root system as a temporary delegation to facilitate multiple different assessment methods are needed now for collecting that data. And right now the discussion group has identified four different methods for collecting and analyzing that data and we'll talk a little bit more about that in a second. Next slide, please.

But this overall risk assessment framework that is being described in the group or the discussion group Study 2 report really has two major goals. And the first goal is ensuring that the name collision data can actually be assessed. As I mentioned, that requires a root zone delegation to gather that data. And this allows the group, the technical review team to be able to look at that data and perform both a qualitative and a quantitative assessment and a risk management kind of process around the potential impact for harms or security threats around that name collision string. And the second goal is then clearly to provide a process for ICANN to evaluate those mitigation or mediation plans. And if needed that there is a mechanism for mitigation and remediation plans to be developed. So next slide, please.

Taking a look at the framework, it's pretty straightforward here. We've identified four different stages. Stage zero is the pre-application period in which applicants are encouraged to look at publicly available data related to name collisions a priori their submission of the application for the next round. There are data sources out there like ICANN's ITHI or the magnitude data sets that provide some kind of insights into the

potential status of name collisions for certain strings. Now while this is helpful for applicants to at least have a little knowledge going in, it in no way guarantees that, you know, the presence or the absence of that string in those data sets guarantees the actual potential risk of those actual name collisions. But here we're trying to provide some insights into, or give some more clarity to applicants before they go in for their submission.

But essentially then we proceed into stage one after the application has been submitted. The technical review team is going to perform a similar step here where they take a look at the publicly available data to make sure it's not a super high risk string before proceeding to delegate that string into the root server system. So once the string is delegated into the root, we're in stage two in which the TRT is now being able to collect the various different types of DNS and other telemetry data to do their risk assessment.

As mentioned before, there are four methods defined in the Study 2 report for collecting data as it relates to name collisions. I won't get into the technical details here, but the first one is no interruption, which really is looking at preserving a no error, no response to the queries that are being sent. This is to simulate essentially the non-existent domain response that collision systems are receiving currently from the root server system. So it's not very hopefully impactful to the name collision systems out there, but it facilitates the data to be collected that's useful for the TRD to do their name collision assessment.

The second method is controlled interruption, and this is similar to what was used in 2012. Here, controlled interruption is both allowing

for that collection of data, but also performing another function, and that is notification to the impacted systems out there. Methods three and four, visible interruption and visible interruption with notification, are a more advanced way of gaining different data related to name collisions that are not looking just at DNS, but are actually looking at data from the impacted systems behind the recursive resolvers, sending those DNS queries. Essentially, at this point, visible interruption and visible interruption and notification require that delegation of the TLD into the root, who is basically had an empty zone, now to contain a routable IP address that will direct all answers and queries for that name collision string to a sinkhole server. That sinkhole server will either be dropping the connections, but being able to log important pieces of information like IPs and ports and whatnot to better understand the underlying cause of those name collisions. Now, once all that data has been collected, the TRT does its magic and performs its risk assessment function on it, and then ultimately, in stage three, submits a recommendation to the ICANN board. Next slide, please.

So, clearly, the TRT is a pretty important role in the name collision assessment framework going forward. The TRT needs to have a deep understanding of DNS and how it works, as well as the internet infrastructure. They need to be able to do data analysis and understand how to put all of that together within a risk management framework. But functionally speaking, the TRT really needs to assess the visibility of name collisions, gather that data, their findings, their recommendations, assess any mitigation or remediation plans if they're needed for a particular string, and then also advise an

emergency response function. Next slide, please. And I think that goes back to Suzanne.

SUZANNE WOOLF: Thanks, Matt. And we did want to highlight in this presentation one specific issue that we know is going to require significant discussion beyond the report and the discussion group. And that's, we've touched on a little bit. You can't just reuse the collision detection methods that we used in the past. It's the internet. It continues to change. So just for instance, there's a lot more IPv6 in the world than there was in 2012. And controlled interruption as implemented then doesn't really work for IPv6. So that's lost visibility. That's lost data. Also for root servers and resolver operators, because of various technical changes and regulatory changes having to do with privacy and security and managing potential risks and so on, there's much less data getting into the open internet than there was. And those changes are very good things in many ways. But they do reduce the data available for an exercise like this.

But to seriously analyze named collisions, there must be, there has to be data from a variety of sources. One thing we discovered is that it's impossible to build a generalized case for root causes. There was going to, there was originally going to be a study three in this investigation and, well Study 2 recommends not doing study three because building a generalized case would not yield incredibly useful results. Also worth noting, impact assessment might require large amounts of data over significant periods of time. And this adds up to, you know, a situation that has to be pretty thoroughly understood. ICANN org in their

comments to the public comment proceeding expressed concerns about risk to privacy and confidentiality with some of the proposed data collection. And, you know, we fully agree that these concerns need to be thoroughly understood and addressed as the DG recommendations move towards implementation. That's an issue that's already been discussed extensively and we expect to hear more and to try to at least respond to the public comments we got when we do the final report.

And I think it's important to note, I'm not sure we did, that the discussion group and SSAC, the work party inside SSAC on name collisions are separate groups and managed according to separate remits. So Ram is going to talk about what's going on inside SSAC about this.

RAM MOHAN: Thank you Suzanne. As Suzanne mentioned, the SSAC itself has a work party and we're looking further into the discussions that have happened, the report that the DG has put out. So if you go to the next slide, on the whole inside of the SSAC and the work party there is strong agreement with the recommendations that have come out from the, and the observations that have come out from the NCAP discussion group. But there are a couple of overarching things for everybody to understand. Measurements at this scale are very difficult. They're not easy to do. And we note that the studies are based on measurements at the A and J root. There used to be quite a lot more data available at the root systems 10, 12 years ago. It's far lesser now. Technology has

evolved and with that evolution the amount of data that is available to analyze has reduced.

As a result of that, much of our own advice, when the SSAC brings out its advice, it's important to qualify that it's not necessarily going to be quantitative in nature. It's going to be more qualitative. And the way to explain that is when we say the risks of name collision are higher or have increased, we do not have a way to say by how much. We do not have a way to say they've increased by 10% or 20% or something like that. But we feel quite confident to say that the risks are persistent and they appear to have increased. So that's just an overall preamble to everything that we're going to be saying. Next slide.

Having said all that, the clear conclusion that we're coming to, just like the DG has come to, is that there is a persistent threat to the security and stability of the domain name system from name collisions. We're quite convinced of that. We think there is unequivocal evidence, data, that points to that. And having said that, the data that we are seeing makes it very clear that it's not enough. The amount of data that has to be collected doesn't compare to what was collected 12 years ago. We believe that not just collection of the data, but when in the process the data is collected and where it is collected are now important considerations to look at.

And in 2012, for those of you who may recall this, in 2012, the way the data was collected was in the name collision framework, name collision assessment and data gathering was done at the end of the evaluation process. So in other words, ICANN evaluated all the applications for a particular string. They decided who was going to get the string and they

granted the string to that applicant. Having granted it, the applicant had to go through a controlled interruption period to determine whether there were problems with that string. And if they qualified past that period, then they were able to actually use the string for their own benefit. We think that with the changes in the technology, with the much lesser amount of data that is coming through to the root servers, we think that that method no longer works. And we also are concerned that the name collision analysis method that was used in 2012, which is controlled interruption, if that same method is applied now, it won't do the job. It's not going to be adequate. One of the factors is the method that was used in 2012, there were far fewer IPv6, there was far less IPv6 traffic and access as compared to now. That's just one thing.

So therefore, where the SSAC work party is moving towards is to be quite clear in saying, name collision analysis should be done as soon as applications come in, not after applications go through a diligence process and when they're ready to be granted. We think that the strings should go through a temporary delegation phase, collect data, and we also think that doing it right at the beginning allows it to be streamlined with the rest of the evaluation process and actually won't delay the strings from going live.

The other thing that has been pointed out in public comments, ICANN.org has provided a comment in a public way that expresses concerns about two of the methods that are being discussed for analyzing name collisions, visible interruption and visible interruption with notifications. And in both cases, the legal analysis or the legal comment is that significant privacy implications apply. We do not have any dispute with that. In fact, we think that privacy implications apply

for just normal controlled interruption as well as visible interruption and visible interruption with notifications. So the privacy implications are universal in all three of the methods.

But our concern is twofold. One is if ICANN decides to first grant the TLD to the applicant and then perform any kind of interruption to analyze name collisions, the very clear privacy risks merely get transferred from risks that ICANN bears to risks that the TLD applicant bears. So you're not actually mitigating or removing the risk. You're transferring the risk and the costs to the applicants. That's one thing.

But the other thing that we're quite concerned about is that a method like that might amplify the risk in a much more significant way. And let me explain that. The reason why we think it might amplify it is because instead of one entity looking at conducting these tests and controlling the privacy implications and the privacy risks, ICANN, which would happen in the method that we're suggesting, if it were to be done and transferred to various applicants, you're depending on two things. One, that all the applicants have the infrastructure, the ability, the capacity to perform appropriate controlled interruption or visible interruption or visible interruption notification kinds of tests. That's one thing.

But the other thing is that they actually will know what to do with the privacy, with the PII that is coming their way and deal with it in a uniform way. So we think actually that the risk vector is greater, potentially amplified in a greater way if it were to be done at the end of the process. So we think in general, so far we've been thinking that don't bolt it on at the end of the evaluation process. You've got a chance to actually do it in a much better way. Do it right when the applications

come in. Perform the evaluations. Provide the data to a well-qualified review team. Build transparent criteria and then make decisions that are informed by data rather than granting a TLD to somebody and then effectively hoping that it's going to be done in a uniform and in a controlled manner. So next slide.

We haven't finalized this, so this is still work in progress. We're sharing it as we come across it. But what we're really getting to saying is that unless there are compelling privacy concerns that absolutely cannot be mitigated, we think that ICANN, the board, our advice goes to the board, we're looking to tell the ICANN board please address the privacy concerns. Find ways to mitigate the privacy concerns. Don't compromise on the critical security and stability aspects of name collision. So that's the direction we're going. On the next slide, we have certain timelines. We are tracking to those timelines. We expect to publish our final report and results to the ICANN board by the 1st of May. So that ends this part of the conversation on name collisions. Instead of going to the next set of slides and potentially having members of the audience forget by the end of the 90 minutes the questions that may have sparked in their heads now, why don't we take questions from the audience on name collisions and also on Tara's presentation on new members and member outreach. So let's pause for a moment. Danielle and Kathy, I'll pass it on to you to manage the queue on this. Thank you. Steve?

STEVE CROCKER: Thank you. So I've sat through a couple of these presentations, but I have not been following this work in detail. Can you say a word of

clarification as to why data collection is more difficult now than it used to be and to the extent that IPv6 is a separate and distinct issue why that is?

RAM MOHAN: Warren, do you want to take that?

WARREN KUMARI: So there are a number of reasons why the data collection is harder. Many of them are related to privacy. Let me quickly find my link. So for example, over the past couple of years, the IETF has added QNAME minimization to the DNS, and that basically means if you're looking up, for example, www.foo.bar.example, you only send the necessary bit to each of the name servers. So the root won't see the full query. It would just see the top part. There's also aggressive NSEC, which is something that's built into DNSSEC and was designed as both a privacy enhancement but also an anti-DDoS mechanism. With that, you don't actually send every unknown query to the parent. You only send one, and then you get a list of names that says nothing exists between here and here. So if .example didn't exist, the root wouldn't see queries for that. It would only see queries for, for example, something between EE and EZ don't exist is the answer you get back.

There's also a local root where a significant number of resolvers are running a copy of the root zone locally to themselves. So these queries also don't hit the root. And I think that that's -- oh, also it looks as though the size of caches have increased over time, which means that

often you will get a response and you will cache that for much longer than you would have before. So the root sees even less queries.

RAM MOHAN: Thank you, Warren. I'll pass to you, Danielle, for any questions that have come remotely. For those who are here in the audience, if you have questions, there's a microphone there. Please line up behind it so we have a visual cue that we can come to you for questions. Over to you, Danielle or Kathy.

KATHY SCHNITT: Thank you, Ram. We have a question from Maxim. When do you expect to have an answer from the ICANN board?

RAM MOHAN: Jim?

JIM GALVIN: Thanks, Maxim, for the question. Jim Galvin, SSAC's liaison to the ICANN board. Although the board is certainly beginning to pay attention to this effort, certainly SSAC and the board had an engagement this morning. I would encourage folks who weren't there at the time, if you want to go back and look at the recordings, it was during the first session this morning, and you can see that there was quite a good set of questions and answers. So the board is beginning to get involved in understanding.

The formality of this is that it won't really be on the board's agenda until after the final work product is delivered on May 1st. And then the board will begin its appropriate deliberations and come to an answer as quickly as possible. Thanks.

RAM MOHAN: Thank you, Jim. Dan?

DAN HALLORAN: Thank you, Ram. And I know many of you, but not all of you. I'm Dan Halloran. I'm with ICANN Legal, ICANN's deputy general counsel. And I'm also ICANN's chief data protection officer. So this is all good, fun stuff for me. And thank you very much for the presentations. I've learned a lot. And thank you for helping dumb it down a bit for us, too, in the slides and the way you presented. I especially like the analogies, you know, the envelope there, addressed to Mary, anywhere, that's the kind of thing that helps us understand that. And so to me, as the chief data protection officer, I get a bit scared of this. Like, whoa, there's these envelopes, and now you guys want ICANN to open up this envelope, addressed to Mary, and see what's inside it. And thank you very much for treating the privacy concerns seriously. And hopefully you've seen -- my associate, Odeline, did a brief memo just with kind of an initial analysis, pointing out some of the issues. And you guys have taken that seriously. And thank you very much. And I just hope to, you know, from here on out, especially step up our collaboration and support for the discussion group, the working party, whatever we can do to help support and provide information if we can. I don't have any questions today of you. But if you guys have questions, I know we

already have a first batch of questions from the discussion group. So just anyway, to say hi and see how we can help support this effort. Thank you very much.

RAM MOHAN: Thank you, Dan. Really appreciate you. Let me go to Peter, who had his hand raised, and then after that to you, Suzanne.

PETER THOMASSEN: It's actually a response to what you just said. So apparently if these envelope analogies help, then I'd like to stick with it. You said it sounds scary if ICANN starts looking into the envelope and stuff like that. Actually the response that gets sent doesn't depend on what's requested within the envelope. Usually you would send an envelope to a web server and say I want this and that, and here's my cookie or login credentials, blah, blah, blah. And then the web server would prepare an individual response based on that. Now what these interventions would mean is that the sender's address on the envelope is used to send a response back, but the envelope's content is not really inspected and there's no personalized answer or anything. It's just using the sender to send something back that makes it visible that it didn't work as expected. But it's not actually looking inside the envelope, except that of course, I mean, it's always there, so technically I could, but it's not the same as actually using it, I would say.

RAM MOHAN: Thanks, Peter. Did you want to respond to that, Dan?

DAN HALLORAN: This is the kind of discussion we need to have a lot more of. So I won't go into that depth. We could probably go on for an hour just on this envelope question. Happy to dig into it more. The issue though is then still under the controlled interruption, as I understand it, that envelope would never get anywhere. You'd put that envelope in the mail and the mail carrier would pick it up and go nope, here, I can't take this, I'm going to give it back to you. And so the mail carrier would have seen it and just handed it back, but actually there's no mail carrier and it just loops right back. Here though now, this envelope is coming into ICANN's office, we're seeing this name, it's going into the memory of our servers or whatever, that's personal data. Mary's name is on there, that's an identifiable human being. We're collecting and processing that data, that creates these data protection issues. And yeah, we can discard it really quick and get rid of it, but there's risks to Mary's personal, her privacy there, and we take that very seriously and it creates legal risks. And so anyway, the goal isn't, this is not, we said I think in Odeline's memo, this is challenging. And maybe even, and just to lay our cards on the table so there's no surprises, it's challenging if not impossible under certain regimes. But we'll sit with you and do whatever we can to see how we can do this, how can we mitigate it, how can we balance. And the important thing to understand too on the other side is what are we really balancing? What is the harm that we're jeopardizing Mary's privacy, what are we trying to prevent here? And we need to look at that. I'm not super familiar, I'm not, I rely on Steve and Matt and others to help me with all that technical stuff, and you guys. But what's the great harm? And I know we have the name collision

reports from last time. I think off the top of my head there were like seven that were categorized as severe harm. And that severe, by severe harm that does not mean a plane flew out of the sky or somebody died on the operating table. That means one company's employees couldn't access the VPN that day until the network admin fixed it or something like that. Or maybe, I'm not trying to minimize it, there were at least seven severe harm reports. But that's the kind of thing we have to, in the end, balance that and figure out if there are ways to mitigate it. But anyway, thank you for the intervention and happy to keep the conversation going. Thank you very much.

RAM MOHAN: Thank you Dan. Suzanne and then Warren.

SUZANNE WOOLF: Dan kind of got ahead of me, but I want to say that I think having Org file a public comment was really helpful to the discussion group and to the community as the conversation continues. And I know that there are challenges even with writing something like that in public, and we appreciate that. Thank you.

RAM MOHAN: Warren?

WARREN KUMARI: Yeah, I'll just second that. Thank you. That was really helpful. I just want to sort of expand on the analogy, and it's probably dangerous because analogies almost always end badly. I think that the thing is more, you

get a letter and you write on it, you know, no such person at this address, return to sender. And so I don't think it's actually you opening the letter, you're just noting this arrived here, send it back and make a note so that the original sender knows that these letters are not getting to their final user.

DAN HALLORAN: Ram, sir, I don't know if you want to... I'm happy to keep engaging. I don't want to suck up the oxygen here or take anyone else's time.

RAM MOHAN: For the SSAC, this is probably the biggest priority thing that we're looking at. So I'm going to devote some more time to it, and we may end up taking one of the other reports that we were going to share with the public. We'll take that for the next time.

DAN HALLORAN: The real issue is just what you said to maybe on a technical level sounds like no big deal. You get it, say no such address, send it back, delete the data, anonymize it, whatever. To a data protection authority, you're collecting and processing personal data, just that step. And there has to be a legal basis to do that. So anyway, this is what we'll work through.

RAM MOHAN: Thank you, Dan. And I feel fairly confident that both Matt and Suzanne will invite and appreciate you, Odalie, and others to come in and be part of the dialogue to help find ways to be responsive to the concerns. We recognize they are real concerns, but we also recognize that we believe

there are real harms. And perhaps we need to do a better job of articulating what those real harms are so that both ICANN org and the board can then come to an informed conclusion on what the risks and the benefits are. Danielle or Kathy, are there any questions online that we ought to address, both on name collision and on membership?

KATHY SCHNITT: We do, Ram. We have a question from Harish. How are we planning to deal with IDN variants? Is it also the part of name collision work?

MATT THOMAS: So, in terms of IDN variants, everything is going to be conducted on a string by string basis. So, there's not coverage for all the variants. You're going to have to do the analysis on a string by string basis, essentially.

SUZANNE WOOLF: I just also want to point out that there have been, I think, multiple PDPs about IDN variants. So, that's covered by those policies.

RAM MOHAN: Yeah, and I was going to say something similar, that there is an LGR framework that ICANN has and variants should be quite well covered in that framework. So, it may not be nearly as important or relevant an issue for name collisions. Any other questions, Kathy?

KATHY SCHNITT: No, but Rod has his hand raised.

RAM MOHAN: Okay.

ROD RASMUSSEN: So, yeah, I actually wanted to go back to continue pounding on the envelope analogy since it seems to be working. So, the things that work, we like. And I would like to add my thanks for the analysis that the legal department did, and it was very helpful to see that. To stretch the analogy even further, I think the area we have more recently been discussing here and surfaced is that it's not the fact that there's going to be envelopes sent out and there are going to be things coming back to folks. It's who's going to be sending them and who's going to be getting them, which in the analysis so far has been done if ICANN had done that. In the previous round, all the TLD applicants did that and under, with no oversight control, etc. So, I think what you're going to hear from us in the discussion group is going to ask for an analysis on the broader ecosystem. So, you know, what are the challenges and opportunities for mischief, etc., if this is pushed out to others to handle basically on their own? And so, what's the responsibilities around for various parties in dealing with that? So, think of it as lots of post offices and lots of different ways of sending envelopes and getting them back and ripping them open or not.

UNIDENTIFIED SPEAKER: Thanks, Rod. And one thing, I think it's a very good point and something to focus on. If we would end up going that way, or I don't know if this is in the realm of thinking, if it does get pushed down to the registries ultimately, it's within ICANN's power and the agreements are through policies to set requirements on that. What exactly would registries,

what must they do, what must they not do? So, it doesn't have to be just kind of open season, willy-nilly chaos out there with this as a possible way. So, again, I think it's going to be really important to sit and talk through each of these things, think it through. Is there a way to mitigate that? Well, here's the risk. How could we address that? Just to try to get to that balancing. Thank you.

ROD RASMUSSEN: And who's the postal inspector?

RAM MOHAN: Thank you. Kathy?

KATHY SCHNITT: Nothing further.

RAM MOHAN: Okay. Anything on membership, Kathy, that is online?

KATHY SCHNITT: No. We've taken care of that in the chat.

RAM MOHAN: Okay. Wonderful. Anything on membership that folks here in the audience want to talk about? Okay. So, I'm just looking at the time. And we have four things. And Barry, I think what I'm going to do is take out SAC 123, that presentation. And I'll note that Barry is presenting SAC 123 at the DNSSEC workshop part three. That's on the –

BARRY LEIBA: I was going to plug that anyway. Part of that. So, yeah, come to the DNSSEC workshop. The last session of it, I'm giving that talk.

RAM MOHAN: Thank you. So, with that said, let's go to the next slide. And I'll hand it to you, Jeff.

JEFF BEDSER: Thanks, Ram. I'm Jeff Bedser. I'm on the leadership team with CleanDNS. This will be quick as more of an advisory of where things are going and what we're kind of watching. So, as most of you are probably aware, the board did approve the change in global amendments to the registrars and registries related to their handling of well-evidenced cases of DNS abuse that they must be mitigated or disrupted. So because we have a paper on that, SAC 115, and we did talk about that whole interoperable approach to addressing DNS abuse handling, it's a topic we're tracking. Again, tracking is probably the best word. We're cautiously optimistic this is going to make an impact in volumes and in reduction of uptimes of the domains used in abusive activities.

So we just have a couple of questions in this active engagement we're doing with the community about this topic. And it's really to kind of see what the ICANN org does at a compliance function with what I would consider the new teeth they've been given to address this issue. And we're looking to see, like, what are the changes anticipated for new obligations? What is next for addressing the other online harms that should be mitigated outside of the ICANN contractual obligations? And

the obligations specifically do call out for disruptions and mitigations, but I don't see it defined clearly in that new language what a disruption is. Mitigation is pretty clear. It's stopped. But what is a disruption? What's an acceptable disruption? And what are the works going to be done to help the contracted parties come to an agreement on what is an appropriate disruption under the new contract agreements?

So these are just issues we'd like to track. We're going to keep on. We're engaging with our outreach to the other constituents, SOs and ACs, to talk about these topics and look forward to, you know, hopefully some really good changes based on these obligation changes.

RAM MOHAN: Thank you, Jeff. Questions for Jeff on this topic? Steve?

STEVE CROCKER: Are there any agreed upon metrics for how much abuse there is in a measurable way? And any idea as to how low you want to drive that and think that you've done something meaningful?

JEFF BEDSER: That's a great question, Steve. As far as measurement of total volumes, there are many different schools. I actually did a presentation on this at Tech Day earlier in the week. I think that the key measurement that I personally would love to see is the uptime measurement, how quickly the abuse is reported and mitigated is the most important factor versus the total number of domains reported. I think what we're looking at primarily is the tip of the iceberg. The sources being used by most

measurement studies are block lists. And I do know that the majority of registries and registrars that do have anti-abuse policies and programs do end up addressing a significant amount of abuse reports that never hit the block list. So we're not seeing them because there's no obligation for them to report them outbound to be counted. So I would assume that based on that understanding that the numbers we're seeing are quantifiably lower than the actual total number of abuses because they have to be reported to get to those block lists. And if people don't report them, then you don't see them. And obviously the correlation is from that.

So overall reduction, I'd like to see a reduction in uptime. When SAC 115 was published, we found in our research and interviewing different registries and registrars that the average time from a report was 96 hours. I think it's gone down significantly lower since then to about 48 hours. But I would love to see it get to the point of sub-couple of minutes between report and mitigation for a well-evidenced report.

RAM MOHAN: Thanks. There's a two-finger from Rod, and then there's a question from the other side of the audience.

ROD RASMUSSEN: And to direct on your point on are there agreed-upon measurements. And I would say agreed upon by whom, number one. There are various reports throughout the industry. I would say that nobody has adopted any of them as the ground truth for what abuse is, even very well-defined categories. And ICANN has used the DAAR as a bit of a

measurement tool. It's now being replaced, as we've heard about this week, eventually be replaced. However, I think this gets into creating standards and policies within the ICANN framework so that Compliance can actually do something to measure against what is actually happening. And there's already been some conversations in various places about what is that going to look like and who gets to decide that. And it sounds like a policy thing that will probably have to go through the GNSO.

RAM MOHAN: Thanks, Rod. Question from you.

GABRIEL APRIL: Hi, this is Gabriel. And this is, I guess, mostly to you, Jeff. I note that within the community discussions about reporting DNS abuse in light of the contractual amendments and the obligation to disrupt, there's often used the term well-evidenced reports of abuse. And I've been thinking a bit about that. And in particular, I'm thinking that not only does the evidence perhaps change depending on the type of abuse or criminal activity that's occurring, but from my anecdotal experience, the amount of evidence or the types of evidence that might be available can very much depend on where in the life cycle of the abusive scheme that you actually get clued in in the first place is. And I wonder to what extent in your conversations with community or your own thoughts on this that evidentiary standards are taken into account. The more proactive we are in catching the abuse, the less evidence might be available.

JEFF BEDSER: Great points, Gabe. I think that the well-evidenced is a direct response to an issue that came up in SAC 115, which was that interoperability issue where the abuse reporting community would say, "This is bad. Do something about it. Here's the domain. Trust us." And the potential commercial liabilities on getting it wrong are all on the party taking the action. So in the case where a well-evidenced case is brought forward that gives you all the information you need to allow you to take a commercial decision to stop a contract, I think that's really where that language came from. And you're right. It does put the burden on now you've told people you need to give a well-evidenced report, what must be contained.

Some of the work that CleanDNS has been doing with NetBeacon, we have started standardizing through some work that Abuse Six has done over the last couple of years for standardizing the type of fields and forms for types of abuse so that when it's reported, you're acquiring certain evidentiary fields must be submitted to demonstrate what you're alleging. And then that can be validated. But the goal is we don't want to put the registries and registrars and hosting companies potentially in the body of problem where they have to have an expert on staff for every type of harm to reinvestigate and validate. Because that's just slowing the friction down, which lets the abuse live longer. And the goal here is to standardize a model where the quicker we can get it to the point where yes, this is the evidence required to take action, the faster the mitigation can happen and less victims can happen on that particular harm.

RAM MOHAN: Thank you, Jeff. Kathy, anything online?

KATHY SCHNITT: Nothing online.

RAM MOHAN: All right. Thank you, Jeff. Let's go to the next presentation, which is Peter and Steve on DNSSEC DS automation.

PETER THOMASSEN: So the DNS historically has been a clear text, non-protected protocol that was very easy to spoof. So when you are in the position between the one asking the question and the one giving the answer, you could just inject fake responses. So DNSSEC was invented to add signatures to responses so that clients can check those. And it has to be set up by domain. So when you register a domain name, then your DNS operator has to, if you want to use that, has to provide the signatures and then you also have to go to your registrar and establish a cryptographic link by changing the configuration for your domain and then DNSSEC is active. That is often very simple. If the registrar also provides your DNS service, then they can just do that themselves on your behalf. If the DNS operator and the registrar are not the same, then this is a configuration step that in most cases so far is not automated except for some ccTLDs.

And so this appears to be an obstacle to DNSSEC adoption. There are studies that show when registrants try to enable this stuff, they get it wrong in about 40% of cases. And some registrars don't provide the

necessary forms for the domain owner to actually configure this stuff. And then if there is a form, the forms are very idiosyncratic, they vary by registrar, there's different input formats, it's all very non-trivial unless you're very tech-savvy.

So some SSAC folk started a work party to address this issue and describe this problem space and what could be done about it. There is a technical specification from the IETF that allows the child DNS operator to put information into the child zone, which the parent can discover, and then use in order to enable DNSSEC and establish the cryptographic link between the child and the parent. So that is based on the parent looking regularly, and so there is some delay until that gets detected and deployed. And there are some other issues, like error reporting, for example, when things go wrong, it's not clear how to do that. So this is just one approach. There might be other approaches of doing automation for this kind of setup so that these manual disadvantages go away, including API-based approaches, or maybe the child DNS operator could send a notification to the parent and then the parent could immediately do the deployment instead of doing it like whatever when the 24-hour timer is over or whatever the interval is. And so we're working on a report that describes this problem space, and we'd like to invite the ICANN community and, you know, ccNSO and GNSO and the IETF community and everybody else who's involved with this to think hard about these things and kind of figure out what's the best way of doing it in order to have registrants enable DNSSEC if they want to use it and remove those deployment barriers.

It also has other advantages, like, for example, if you do things manually, there's always room for error. When you do things

automatically, assuming that you don't have many bugs, so no bugs have been encountered and fixed, then that decreases the margin for error. So the hope is that DNSSEC will, by adding more automation to it, become more stable, less brittle. And yeah, so that's essentially what we're working on. And there have been quite some changes to the document. There was an almost final version last year in fall, then we realized, oh, maybe that was a little, like, one-sided in a few aspects, so we're rewriting some parts of it. So in case you've heard about the document being closed to publication before, that's changed, and we're kind of redoing some of it, and then I guess we'll publish it sometime, I don't know, over the next few months. I'm happy to take questions.

RAM MOHAN: Thank you, Peter. Questions for Peter and Steve? Kathy, anything online?

KATHY SCHNITT: I have a comment from Maxim. For DS record automation, please consider talking to tech ops of CPH, where tech persons of ROs and RRs discuss things and comment.

PETER THOMASSEN: That's a very good note, and I'll take note of it.

KATHY SCHNITT: Nothing further, Ram.

RAM MOHAN: Thank you, Kathy. Any questions from here in the room? Okay. Gautam, it's your turn.

GAUTAM AKIWATE: Thank you. Okay. So the registrar name server management work party is essentially a result of unintended consequences. SSAC has been looking at issues where we look at expired, like the risks that expired domains create for a very long time, and initially, like maybe ten years ago, these manifested in the way of orphan glue records. SSAC had put out a bunch of reports saying why orphan glue records are bad, and then there was effort to fix those, and in an attempt to fix those, I think we have had some unintended consequences where the under-specification of how things should be done has led some registrars to essentially do things in a slightly unsafe way, which is basically what this work party is studying. And this work was started as a consequence of academic paper that me and KC Claffy had worked on where we found that in the last year, nearly half a million domains had been implicitly exposed to domain resolution hijacks, and so part of this work party is trying to figure out -- next slide, if we can -- part of this work party is trying to identify practices, sort of trying to fix this issue once and for all, but as we know, like I think we're going to have to keep monitoring it, but the goal of the work party is to give options for registrars to implement that we believe would be safe and would prevent exposure for domains. So essentially what the end work product would look like is for each of these options, what the benefits, burdens, and the residual risk for each of the parties, and the parties

that we are sort of looking for are the registrars, registries, and the registrants, and yeah, I think we are -- we have a work party product that's out for review, and hopefully by next ICANN, we should have something that is more -- that we can share with the wider audience.

RAM MOHAN: Thank you, Gautam. Questions for Gautam? Any questions from the room? Any questions online, Kathy?

KATHY SCHNITT: No.

RAM MOHAN: Okay, Gautam, you got away with this, so -- this meeting, anyway. Okay, so that brings us to the end of this session. Just in summary, if you go to the next slide, what I wanted to do was just briefly ask each of the folks who presented take a minute or less than a minute, 20 seconds, to just, you know, provide your view on what you were talking about, just to refresh the audience's memory on the various topics that their SSAC has been looking at. So first, let's go to you, Tara.

TARA WHALEN: Sure thing. I guess to recap, we are open for membership. We are eager and excited about welcoming new members. We're going to do everything in our power to let you know how you can join, what you might need to know in order to join, and to set you up to work among us. And we'd be excited to have you among our ranks. And please come

to our sessions, look at our papers. You know, we want you to meet us. We're a friendly bunch to work with, I promise.

RAM MOHAN: Thank you, Tara. Suzanne?

SUZANNE WOOLF: Only that there's more work to be done. And if you want to be part of it, please join us.

RAM MOHAN: Thank you. Jeff?

JEFF BEDSER: I'll take a page from Ram Mohan's book. I'm going to say not only how do you measure success, but how do you measure so we can measure success?

RAM MOHAN: Thank you. Let's go to you, Peter.

PETER THOMASSEN: So of course, the DNSSEC folks are the most fun to work with. So if you want to join, that's cool. And yeah, thank you, Jeff, for the measurement perspective.

RAM MOHAN: And Gautam?

GAUTAM AKIWATE: I guess part of what we realize in the work party is constant vigilance, which is whenever we think that we have solved an issue, there are new issues that crop up as a result of unintended consequences. And this has been a lot of fun to work with. A lot of SSAC folks have shown up. So if you would like to be involved, I think we're nearly at the end, though, hopefully. Looking at Peter's journey, I am hopeful that ours is slightly less turbulent.

RAM MOHAN: Thanks Gautam. And Barry, do you want to give one last shout out for when you're going to be presenting?

BARRY LEIBA: Okay. I was thinking of shouting something else out. This afternoon is the DNS Sec and Security Workshop. There are three sessions this afternoon. I'll be presenting at the end of the third session. So please come to all of them. And Jacques, please say more, because Jacques runs the thing.

JACQUES LATOUR: Yes. So 1:15, we have DNSSEC and Security Workshop. The first panel is on digital trust. We have a lot of content. It should be very interesting. Try to be here on time, because we're starting, and we have a packed agenda with really cool stuff.

BARRY LEIBA: Jacques likes to crack the whip on the time. But the thing I wanted to say is about the membership stuff. Personally my key diversity point is getting more younger people earlier in the career, people with different views of the Internet than those of us who used punch cards back in the day and ran around with sheets of paper with the routing tables on them. So it's a different world, and we want different views.

RAM MOHAN: Thank you. On that note, we conclude this session. Thank you, everyone, for coming, and thanks for the SSAC members for being here and being part of the journey. Thank you so much. We are adjourned.

[END OF TRANSCRIPTION]