ICANN80 | PF – Joint Session: ALAC and SSAC
Tuesday, June 11, 2024 – 15:30 to 17:00 KGL

YESIM SAGLAM: Hello, and welcome to the joint meeting of the ALAC and SSAC. My name is Yesim Salam, and I'm the remote participation manager 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 will only be read aloud if submitted within the Q&A pod. If you would like to speak during this session, please raise your hand in Zoom, and when called upon, please unmute to speak.

On-site participants will use a physical microphone to speak. Interpretation for this session will include English, French, and Spanish. Please state your name for the record and the language you will speak if speaking a language other than English. Please remember to speak at a reasonable pace. With that, I will hand the floor over to Jonathan Zuck, ALAC Chair, and Ram Mohan, SSAC chair.

JONATHAN ZUCK: I guess I'll go first since this is our room. Some host privileges here. We need jerseys or something like that. The visiting team is the SSAC. Welcome, everyone. This has become a regular meeting between the At-Large community and SSAC. Because we have a lot of shared interests in the fate, if you will, of individual end users and their interaction with the DNS on the internet, and making sure that policies that get made inside of ICANN really take into consideration some of

those downstream consequences. I think we've had a really productive partnership.

We've been able to mirror and amplify some of the messages that have come out of, of SSAC. As we'll talk about in a minute, the SSAC is one of our potential clients, if you will, for our new campaign journey that we are on through the Safer Cyber campaign. So, I'm excited to have you here. Welcome to the SSAC. With that, I'll pass the microphone to Ram.

RAM MOHAN: Thank you. I don't think the Jersey idea will work, because that signifies opposing teams, and we're not that. It's such a pleasure to be here. There's a whole bunch of other members from the SSAC here in the room as well. If you'll briefly raise your hands so folks here know that there's a--

JONATHAN ZUCK: That's too many.

BARRY LEIBA: We're everywhere.

RAM MOHAN: It's a flanking move. Pleasure to be here. We have a really good set of topics to discuss. I think there are some evergreen topics that will always stay between our two groups and these areas of interest. I'm really pleased that we have this now as a regular feature, because there is a lot to go and do, and to build a better and a safer environment for

those who use this tool called the internet. With that, back to you. Jonathan, I think you're going to start with the framing of the Safer Cyber campaign that we're looking to do, and then we have some specific content that I think will be very useful for the membership.

JONATHAN ZUCK: That's excellent. Thank you. Just a reminder to staff, I need to be a co-host here, at some point. So, we've had different topics of discussion. This is talking to my fellow members of the At-Large community about trying to organize ourselves around this idea of campaigns. Because more and more different communities within the ICANN broader community are coming to the At-Large community to ask for help in amplifying a message out to the broader populace or to try to get feedback on a particular issue that we're discussing from a broader number of people.

The third possibility is education. What can we do to educate end users on a particular topic? And that's the issue here, that the SSAC has come to us with an idea of a Safer Cyber campaign about reaching out through the RALOs and ALSs and teaching a course, much like the UA course from UA Day, but focused on people protecting themselves online. We've begun the creation of this course.

The idea is to create something very much like they did with UA Day. I've been calling it a seminar in a box, "Here's your presentation, you can modify it to personalize it for the presenter. Here's a video of somebody delivering the presentation. Here's some Q&A. Here's a few hundred dollars for snacks and things like that." The whole idea, like UA Day is,

"Here's something ready for you to deliver to your constituency, to your group."

And so, we've been working a little bit on what this one might look like. This is the first draft. It's still a work in progress, but just wanted to give you an idea. We're trying to make this kind of slick, and something more like what you might see Apple deliver as opposed to something you might see ICANN deliver. Oh, I didn't say that, I'm sorry. We're trying to make it very user-friendly as we go along. This, when it's done, will probably be a two-hour course.

I'm not going to take two hours to go through it, but just wanted to give you a little sense of what we're doing to try and put this together. There's animations built into it and there's data. There was a global cybercrime in-- I don't know why that says 2021. That should say 2023, I think, or something. There's still lots of errors. $6 trillion. Basically, it's been going up 50% a year for the last number of years. And then it's $10.5 trillion in 2024.

What we're going to do is talk about different types of cybercrime, hacking, financial fraud, phishing, online scams, identity theft, et cetera. But where we're putting a lot of our focus is on this issue of phishing, because that's where the end user is most involved. They receive an email that says, "Hey, call this number here, click this link here, download this attachment..." and that's when the trouble begins. And so, that's what we want to start to teach them about. We talk about phishing. We're going to talk about what phishing is, how to identify it, and how to prevent and report it.

Again, we've got some animations, some graphics for what this looks like. We're going to talk about what phishing is. This is different types of phishing: spear phishing, whaling, smishing, et cetera. Then we'll talk about how you identify phishing. What are some of the common tactics and warning signs? Is it a suspicious sender? Is it urgent language? Is it a request for personal information? Things like that. Those are your warning signs.

Then we're going to go through various sample emails and show these different warning signs in action, from actual emails that have come through. These are just a few that I had lying around, but we'll come up with more to understand what kind of language is used. AI has made this even harder to identify because they pull personal information from you on the web and customize the messages that you receive and things like that. When you first get an email, you're asked to do one of three things.

One is to download an attachment, one is to link to a pharming site, and then the other is to call a toll-free number. Those are the big three that get asked. We're going to talk about different types of malware so you just get an idea of what this will look like and things to avoid with attachments, things to avoid with 1-800 numbers, how they impersonate customer service personnel and things like that, and what they request. I know I'm going through this quickly, but this will be something that -- this is a link on the agenda, so you can go through it. And we're going to spend a lot of time making this as good as it can be, but this is the first draft of it.

Some tips, for example, the hover method. When you hover over a link, it tells you where you're going. But then here's some examples of some homophones that don't work in that method, because they're using IDNs, for example. Just going through different examples. Then there's a quiz that Eduardo found at one point, that Google put together in multiple languages, that helps you identify phishing emails and things like that. We had people take that. We'll talk about different types of phishing. Then there's some more graphics coming. Then eventually, we'll talk about best practices. And then finally, We will talk about how to report it. The first step is how to recognize it and avoid it yourself, and what your own hygiene should be to perfect it.

The other thing, again, that Eduardo brought to our attention are the tools that have been created inside the community to facilitate reporting it, including that beacon and the asset tool. We're going to go through and show how to use those tools as well, as a user, if you're willing to go that extra step of reporting it and trying to help the rest of the community out as well, once you've avoided it.

That's basically the structure of the course. We're in communication with the SSAC to improve on that and to add things to it, et cetera. But I just wanted to give you a little bit of a preview of what we're trying to build, and that we will then translate into multiple languages and try to put together in a package the way that they did for UA Day. That's the idea. Any questions on that quickly before we move on? Yeah, go ahead.

EDUARDO DIAZ: I'm sorry, I came late. I didn't see the beginning. But I have given courses like that. The first part of what we do is we explain how you read the domains, they are read from left to right. Then I go into the schema that is in the URL so they know the parts of it. Because when they get this humongous link, they don't know what it means. I also explain how the DNS translation goes around, what happens when you hit that key. You can put all this stuff in context. When you show the things, you can say, "If you look at this schema, this is what's happening..." Just as a comment.

JONATHAN ZUCK: Thank you. I'm not going to absorb all of that right now. We should make sure that that's part of a process of improving the -- but some of that stuff is already in there, but help me make sure that it all is for sure. I'm looking for higher-level questions or comments at this point because I won't absorb something that's specific.

WARREN KUMARI: This looks like a really good presentation and much better than most of the anti-phishing things. The big problem is, how are you going to get people to see this? Like, how do you distribute it and get the people who need to watch it to actually watch it?

JONATHAN ZUCK: All of these are experiments. One relatively successful experiment that was done recently by the Universal Acceptance Steering Group was to put together a seminar in a very easy way for people to deliver it, then

they accepted applications for people that have audiences to which they would deliver it. To those people, they awarded them a couple hundred dollars to pay for the snacks or something like that and then delivered this seminar in a box, basically to them. We, very loosely, participated in that process and ended up delivering half of them, or something.

We have within our infrastructure both regional organizations and even local organizations that are members of those regional organizations. For example, a regional At-Large organization could have among its membership a local ISOC chapter, and they are giving presentations every month, and they're looking for things to present onto their membership. The idea then is to take something like this that's packaged up so easily to deliver that's like, "Here. All you need to do is fill the room, and we'll pay for the snacks. And here's the presentation." That's the idea.

WARREN KUMARI: Great answer. This is why the SSAC/ALAC collaboration, I think, is going to be very fruitful. Because safer cyber is something that is of relevance for the people of the world and for users of the internet. We have a lot of content. We have a lot of information that we can provide, and then you can package it and then present it.

MATTHIAS HUDOBNIK: Just thinking out loud, an option could also be to record the session and upload it on the ICANN Learn platform, and then people could access it or you could refer to it.

JONATHAN ZUCK: Those are all great suggestions. We're trying to work through our infrastructure at At-Large, because ICANN Org wants us to use that infrastructure to evangelize the new gTLD program, for example. Others may want us to figure out messaging for DNSSEC, as we tried to talk about at one point with Steve until he went off on another path. The idea is if we can really decalcify our own communications pathways with this, then they can be used for multiple campaigns. But once the materials are created, and Matthias, I think you're exactly right, anybody could give it.

In other words, I could get invited to speak at a conference, for example, and then this is ready. It's in a box, I can take it and deliver it as my content at that conference. It could be on ICANN Learn. It could just be a video, for example, that's on the At-Large or the ICANN website. We also have the beginnings of a partnership. I forgot whose idea it was, one of the Canadians, that we should partner with libraries.

Well, I've started a relationship with the International Library Association. Libraries have become community centers. They are looking for content and things to teach in their community centers and things. They're very excited about having a phishing course that they could take out of a box and deliver. All those things are possible, Matthias, and great examples. Once we work out the bugs ourselves, then it becomes this product that's used elsewhere.

WARREN KUMARI: I forgot to put my hand down but seeing as I have the talking stick again, this seems like one of the best things out of ALAC that I've seen in a very long time. This seems like it could be really powerful and really useful, and I just want to say, wow, this looks good.

JONATHAN ZUCK: Thanks, we appreciate it. We spend a lot of our time here at ICANN meetings complaining to the contracted parties. "Why aren't you doing more to protect users from malicious registrations and things like that?" And we don't have any plans to yell less about that. But I think part and parcel to that is to take up the other side of that equation and say, "Hey, when somebody calls you on the phone and asks for your password, maybe, don't give it." If we can address the problem from both sides, then we're really serving the purpose of the interest of end users.

SIVASUBRAMANIAN MUTHUSAMY: It looks very powerful. Through ICANN Learn and through seminars, you're, again, targeting people who already have some technical background. Shouldn't there be a lighter YouTube version, a non-interactive version, and shouldn't ALAC be aiming at half a billion hits or a billion hits, by whatever means?

JONATHAN ZUCK: That's a great suggestion. I think what we're trying to do right now is focus on what we believe to be the membership of our ALSs. That's our infrastructure, that's our membership. And what is their audience, and

how do we get to that audience? I think that's what we're focused on. But I think as we get some success with this, and success is recognized by the greater ICANN community, we can do all kinds of creative things. As I said, if we do something with libraries, that's a perfect example where that might be an instance where doing something more aimed at kids could be a good idea. Thanks for bringing it up. Sebastian, go ahead.

SÉBASTIEN BACHOLLET: Thank you very much, Jonathan. Just to be sure, it will be translated in different languages, but I guess we may wish not to stick with the ICANN-supported languages because we are much broader. We work in those six languages (or seven) but for this type of document, it must be useful to have a lot more languages, if it's possible. Therefore, how we can organize that, how we can finance that, will be a good topic for us in the future. But, thank you very much.

JONATHAN ZUCK: I think that's an excellent topic. I've been working with some designers -- which is a lot of fun -- on the slides, so they don't just look like bullets and things like that. Part of what we're trying to do is really make it graphical in nature and then put a script with it. A script is easier to get translated than having lots of text around all these slides. It's the script that will need lots more translation than the slides themselves. That's the hope. That's what we're trying to do. I hope we can organize to make it available in more languages. The nice thing about that Google quiz -- you should all try it -- that is already in about 20 languages, or something like that, right now. I have a caller.

BUKOLA ORONTI: During the presentation, I noticed a clause that said "another way to detect phishing is using AI-ish words." I just want to point out that some words that are common vocabularies in some of the African countries that sound AI-ish may not be enough to detect phishing. In such cases, it may not be enough to detect this kind of phishing.

Based on experience, words we usually use are incorrect spelling of English words, talking now from the Nigerian perspective. There are some words that seem to be AI-ish but they are common in our vocabulary. I think we might want to note that, because "sounding Ai-ish" may not be enough to detect because they are common words for us. But then, incorrect spelling or incorrect use of English words is something that we commonly use to detect spam or phishing. Thank you.

JONATHAN ZUCK: Again, I want to encourage you not to get too specific right now because we'll have a process to go through and improve the process. But briefly speaking, what that slide was saying is that AI has made it easier for the criminals. Because they're making fewer errors now in the language. They're doing research about you from what's on social media and things and constructing an email that's specific to you and doing more spearfishing. That's what that was about, is that AI is working against us, not that we can rely on it. This is all about not relying on anybody but yourself to detect whether or not something is a bad email. That's the whole idea of the course.

ABDELDJALIL BACHAR BONG: I think that the presentation is very well done, very well made. End users do need this type of presentation, and it's going to be ideal for the ALSs as well. We can use this presentation for the ALSs in the local internet governance schools and the DNS forums. That can be very useful as well as a document for the end users. We noted in the chat that in social media, we have lots of issues with phishing and other issues, and it's going to be very useful. You talked about figures for 2024. It was $10 billion. I don't know where your figures were coming from, but it was already $10 billion, I think, for 2024.

JONATHAN ZUCK: That's right, some study. I don't remember either. It'll be documented though as part of the final presentation. I want to leave room for the new topics coming out of the SSAC.

PARI ESFANDIARI: I think there may be a need to adjust it for various regions in the long term. At the same time, I think that this is a great tool to use in order to energize the ALSs. That could help by giving them content, expand the message, and bring more of them, energize them. It could work in two ways.

JONATHAN ZUCK: I think that's a very good observation, and we need to work out our own processes. But I anticipate that one of those steps in the process is for the RALOs, for example, to look at the presentation and see how it might

want to change for that region or something like that. Because it may mean bringing in a different sample email or something like that. I think customization is a good idea. That should be one of the steps before they then push it to the ALSs to do that region check or something like that on the materials as we go along. We'll figure out our process. With that, I will hand it back to Ram.

RAM MOHAN: Thanks, Jonathan. On the agenda, we had discussion on disseminating SAC074. But you already get the sense of what we're trying to do. You will see some material that's being presented that's going to be somewhat technical, detailed, et cetera. The intent is to take material like that and eventually convert it into something that is campaign-ready and just digestible. You can cross that item off the agenda when we -- we had to get there.

Let's get to Merike. A little while ago, the SSAC published SAC074. We wanted to bring that forward to you and to speak a little bit about-- We think it's got significant relevance to end users, to registrants and we thought there's no better person to help present that than Merike. Let's get that slide up. Merike is online and should be able to speak to it.

MERIKE KAEO: Great, thank you very much, Ram. Greetings from Japan. I had a conflict at another conference but we've got the internet, which is a great thing. So, why an advisory on credential management? We started the work in 2013 and it was completed in 2015. We already had two existing advisories, SAC040 and SAC044, that did have

recommendations on strong identification and authentication, but they weren't very specific in terms of methods and methodology. In the Durbin ICANN meeting in 2013, at the Tech Day, Maarten Van Horenbeeck gave a talk about DNS hijacking.

He was working at Google at the time, and over a four-year span, he had shown and talked a little bit about 30 different countries where Google had issues with some kind of a domain name compromise. Out of the three top issues that he listed, one was related to credential management and that the attacker attempts to authenticate using a list of frequent passwords or using passwords stolen from another registry authentication database.

Following that Durbin meeting, a number of us in the SSAC talked about, well, credential compromise is still a big issue, should we do something about it? And there was enough interest that we created this work party.

The next two slides are going to be eye charts. They are tables that are taken directly from SAC074. The only reason I'm showing them is because we started to list the different types of credentials that were being utilized in the DNS ecosystem, what they were used for, and also which entity was being authenticated (who's using the credential), and then which entity was validating the credential. The entity being either the registrant, the registry, the reseller, the registrar, proxy privacy, service provider, what have you.

If you go to the next slide, it just lists some more of these credentials. What we're talking about here is really the credentials that are relevant from a registrant perspective, so the relationship that they have with an

MDNS provider, with a reseller, or a registrar. Those are really relevant for the registrant. If we're going to be focusing on that with work that we may do with ALAC, or if ALAC wants to take some of this work and utilize it to disseminate good credential management practices, that's the specific focus we have here.

One of the things that we also point out is the multiple ways that credentials get compromised. And this is usually interesting for some end users because they're like, "Oh, look at all these different ways." And it's sometimes an eye-opener in terms of, no wonder there's so much credential compromise that exists. I'm not going to go in detail in any of these just because I want to make sure I don't run over in time, but you'll have access to all of these slides. SAC074 also looks at all of the stages and introduces the credential management life cycle framework.

Basically, if you're creating a credential, changing it, or renewing it, you have to think about all of these different steps. How do you distribute it, how do you store it? Do you need to encrypt it? How do you recover a lost credential? Delegating. Let's say a registrant wants to change registrars, you're basically delegating some of the authentication aspects, provoking a credential, or destroying it.

So, SAC074 goes into detail into all of these stages, what they mean gives a context. It also talks about best practices. Now, the SSAC didn't make up new best practices. It also points to the existing standards that are utilized by the industry by common criteria or the ISO 27000 certification-related requirements that also point to best practices. If you're using credit cards, there's a PCI DSS standard. We give some

recommendations, but we also, of course, point to industry best practices.

I think we all know at this point how important two-factor authentication is. The document also talks about some of the methods, some of the stronger two-factor authentication methods, some of the weaker ones that can be more easily exploited. One aspect that I always hear about, and have for about 15 years when I give some security talks myself globally, is that people always say, "You know, the experts always tell us what to do, but not how." The how is usually missing.

This link is from a presentation that was given from a registry in Brazil at a LACTLD meeting that I attended where they were presenting, because we were doing some research about two-factor authentication. "What should SSAC say?" "We've got to take different geographic regions into account." Wouldn't you know it, they had an excellent slide about how to do time-based one-time passwords or HMAC-based one-time passwords. The good, the bad, and the ugly, so that others can learn from what went wrong and then how they fix things.

As a synopsis for SAC074, basically, its intent was to provide best practice guidelines to registries, registrars, and registrants in protecting the credentials. Here we're mainly focusing on the registrants. As I mentioned in the slides before, it explores the set of credentials that a registrant manages for registry and registrar, but we want to focus here on the registrant. It deduces the credential management lifecycle framework and then how the framework can be applied to protect registrant credentials.

Recommendation number four is the one that's really relevant to this conversation, and you'll hear from OCTO following myself. The recommendation reads that the ICANN Board should direct ICANN staff to facilitate global hands-on training programs for, in this case, registrars and registries based on the best practices outlined in the document with the goal to enable parties to learn practical operational practices for preserving security stability of the credential management life cycle. So, in a nutshell, SAC074.

RAM MOHAN: Thank you, Merike. Questions?

SATISH BABU: First of all, thanks for an excellent document in SAC074. For us in At-Large, SAC documents have always been a source for authoritative information. Second, the kind of training material that Jonathan showed us was targeted to the typical end user. For that reason, because we're talking about end users who are maybe not technical people, the material has been dumbed down a bit. But we also have our ALSs, many of which are technical organizations, like ISOC chapters. Now, just like UA Day material was meant for technical people as well, like for developers, this material also has a lot of utility in addressing that part of our At-Large structures.

JONATHAN ZUCK: I'll just say to follow up with Satish, I think the aspect of this that most overlaps with us is the registrants as opposed to the registries and

registrars. We certainly consider registrants to be end users, but a particular class of end users that represent a kind of minority of the end user community. But there are certainly, as Satish pointed out, within our ranks, many, many people that own domain names. Getting this information out through our network, I think, would be a good idea, especially the how-to piece of it. Because Merike's right about that. We're told what to do, but not always how to do it.

YAZID AKANHO: Jonathan and Merike have already set the context so well. I can go straight to the point. Technical engagement has the mission to conduct the capacity development, capacity-building activities around the ecosystem to the different stakeholders within the community. Of course, we do develop different course materials around our courses. Currently, we provide the courses in three languages, mainly English, French, and Spanish.

We have a regionalized team. Each of us focuses mainly in the respective regions. We have conducted a lot of training. This is our core business. We do develop training materials, presentations. We run workshops. We also run, of course, hands-on the registries, the registrars, the DNS service providers, in general. But we do also conduct capacity development with stakeholders such as law enforcement agencies like universities, even regulators and government people as well.

When it comes to the credential management we are talking about here, we have a course which is fully dedicated to it. You can go to our course catalog on the ICANN website. Course number 15 is that one,

but recently we have started revisiting all our courses because those courses have been around almost four years now. We have introduced the Credential Management Lifecycle aspect into six of our other courses that we have identified as relevant to also discuss. When you basically take the registry operations for country code top-level domain registry operators, for instance, or when we talk about the DNS abuse rights, it is obvious that discussing credential management is critical.

The Credential Management Lifecycle course itself introduced the concept of "credential." What it is, what are the different versions, if I could say, of credentials. Where are they used in the overall DNS ecosystem? We all know that the DNS is really distributed. There are multiple systems that interact with each other that are managed by different entities, by different organizations.

The audience is primarily the engineers who oversee the operations of those DNS systems and infrastructures. But we do also engage with policy personnel, as I mentioned before, those who write the policies. But also, the managers. Because they are, most of the time, in charge of developing the procedures, the policies, and making sure that the engineers and the team who operate the different systems actually follow those policies and those practices.

In this course, there is also a need for awareness toward the managers. The duration really varies on the type of engagement we are conducting, but basically if it is an online course, it is about 90 minutes to 2 hours with all kinds of interactions, questions, et cetera, et cetera. But we also conduct this in our hands-on training. Sometimes when we do hands-on capacity development with different stakeholders, we also

make them discover how to set multi-factor authentication, for instance.

Sometimes people do think that they need to buy some expensive systems to set this but, basically, even using your native Linux system, you can implement that. So, those are practical cases that we implement during our courses. Of course, the DNS 101 is a prerequisite course to the Credential Management Lifecycle course, and the course outline touches on various aspects of what Merike has already explained in her previous presentation.

This is briefly just to have an overview of our online lab material. It's a virtual environment with provision, and each participant has a set of virtual machines that run on the Linux Ubuntu. This is how we conduct our training, our capacity development hands-on activities. Recently, I think last week, I was in Cameroon with the ccTLD, the registrars, and some system administrators (those who oversee the DNS servers and infrastructure over there). We discussed the Credential Management course, and the same also is done in Latin America and in Asia Pacific and also in Europe regions with the different stakeholders over there.

So we are probably the right resource or the right people who can really help disseminate the Credential Management Lifecycle principles and best practices within the ecosystem, and we'll really be happy to also engage, as Satish mentioned. We do run the training on capacity development as well. I think there are multiple avenues that can help all of us to help disseminate much more of this within the community. Thank you, and back to you, Jonathan.

RAM MOHAN: We've already provided a preview of what we could do. Satish had the useful addition that perhaps there might be two tracks, the end-end user track and the tech-end user track, and we might have different level of material for that. So I think that's a good takeaway. You have a hand up, Claire?

CLAIRE CRAIG: I just wanted to say, I looked at this presentation and I thought it was very informative. Even as we talk at ALAC, at the, as you said, the end user level, we do interact with some of our technical communities in different areas. From where I sit, we do a lot of work through the office of the CTO and have people come to some of the sessions, but I wasn't aware that you had a catalog of courses from which we can now select. This is really informative and can help us in some of the sessions that we plan to have moving forward. Thank you so much for this presentation.

RAM MOHAN: Thank you, and a lot of thanks goes to Merike for presenting this all the way from Japan. Before we go to the next topic, which is on name collisions, one thing that I wanted to share with you is, inside the SSAC, we now have some level of consensus that we have a few topics on which we have a persistent level of interest and information that we expect to stay engaged in. Those are, in no particular order, DNS abuse, new gTLDs, but the SSR (the security stability aspects of new gTLDs), DNS protocol issues, DNSSEC, alternative namespaces, and finally, SSR aspects of internet governance.

Those are six areas where we expect to develop material, content that collates, curates information from what we've done already before, and what we plan to do going forward. My thinking is that we get those areas together that might form a pretty fertile ground for you to pick and choose and say, "There's interest in this area, let's go grab content and then build it out for the communities that you want to serve."

YAZID AKANHO: I forgot to mention that the community can request some training. So drop an email to OCT at www.icann.org, and we will be happy to engage with you and plan for the engagement. As I said, we have staff in all the regions. Myself and my colleague Simon Mayoye, we oversee Africa and the Middle East, but we do have Champika for Asia Pacific. We have Nicolas for Latin America. We have David Huberman for North America, and Ulrich for Europe. Thank you.

JONATHAN ZUCK: We forget, sometimes, all the things that OCTO does, so it's good to get the reminder.

RAM MOHAN: Okay, great. Thank you so much. I think it's time to move to the next topic, which is Name Collisions. We have produced some advice on name collisions. It's an important topic. Matt, I think you're online. Let's hand this over to you, Matt, to kick us off.

MATT THOMAS: Thanks, Ram. So, I'll be a little bit brief and I apologize for any wisp in my voice. I've been battling some kind of cold for weeks. But, happy to talk a little bit about name collisions in our work that we produced that took us several years. In general, I want to just lead off with, name collisions are a continuing issue since 2012 when SSAC started to publish their advice around this.

Name collisions can introduce network disruptions, they can expose sensitive data, they can create attacks. They are available for multiple disruptions or attack vectors and exploit vulnerabilities. Now, as we have seen over the last decade, new technologies have evolved and they have introduced other name collision scenarios. As we go into the next round of new TLDs being delegated, the introduction of more new TLDs directly increases the likelihood of potential collisions.

Now, from the past perspective of 2012 to now, there are several different major notable events. One of those is that the reliance on root server data alone is not sufficient as it was in 2012. There have been numerous evolutions of the DNS protocol as well as technology changes that have impaired the telemetry that is observed at the root server data that was traditionally used in 2012 to observe and assess name collision risks. That is just not the state in which it is now.

The other thing is, looking at the past how DNS data was looked at and collected from a centralized repository since a DNS-OARC and the day in the life of internet data is no longer really that much of a feasible option, which leads to only looking at a subset of telemetry data gathered at sparse sets of root servers in which you are going to get inconsistent and incomplete protections or analysis of potential

collision strings. In general, this leads us to a situation in which we have a lack of consistency as there is no centralized ability to observe all of the telemetry related to a particular name collision string, and it relates to the next round because it influences how the guidance of name collisions should be done in an ongoing sustainable and repeatable manner.

The NCAP Study 2 or the NCAP, being the Name Collision Assessment Project, really focuses on establishing a framework to create a sustainable and repeatable and deterministic framework to allow for the facilitation of name collisions on an ongoing basis. And this includes four key features. The first of which is Integrated Risk Assessment. If you recall back to 2012, name collisions were somewhat of a bolted-on item as to the back of the name collision process and the granting process back then. Here we are embedding name collisions into the broader review process and including it in the overall process not to make security as a bolt-on at the end.

Second, is the establishment of a Technical Review Team, or a TRT, which really is the introduction of a new team to evaluate the proposed new gTLDs based off of the telemetry data that is gathered from the various points in the DNS hierarchy for their empirical analysis.

Third, is the introduction or the proposal of new Enhanced Data Collection mechanisms. These are new mechanisms that allow for the collection of both quantitative and qualitative data from both publicly available datasets as well as those that are tailored to the specifics of an applied-for TLD to enhance the name collision risk assessment. And

fourth, is the prescription of allowing for multiple different assessment methods in terms of collecting that data.

The name collision project really came up with a robust set of different methodologies weighing both privacy versus the verboseness of the data to allow for how to assess name collisions in a varied set of ways. In a straightforward process, the name collision process is, we'll say, four stages. Stage 0 begins with the applicants looking to identify potential issues before they submit their application to ICANN by looking at publicly available data, such as ICANN's ITHI magnitude data or some of L-Root's data that presents some leading indicators as to potential risks around name collisions for strings that they may be applying for.

In no way is that an indicator of acceptance or denial, but it at least gives some a priori knowledge to the applicant that there are or are not potential name collision risks. I'll qualify that with, even if their string does not appear on that list, that doesn't mean entirely that name is not void of name collision risks. There is an effort to at least give some kind of information to the applicant a priori.

Once the applicant has submitted their application for a string and it goes into Stage 1, the Technical Review Team will look at that string as well and look at other available data sets that are available to them, including some of the public ones that were available to the applicant pre-application, in which they will look and make a decision of, is this a string of high risk or if it's not high risk. Once that is gone on, if it's not a high risk, it will go into Stage 2, into the name collisions, where ICANN will temporarily delegate the TLD into the root zone.

Then the Technical Review Team will conduct one or more of the following assessments: no interruption, controlled interruption, visible interruption, invisible interruption, and notification. Those were some of the multitude of the techniques that I mentioned before. They will clearly be tailored upon the application string. They will be tied to the discretion of the Technical Review Team. This is a method in which the Technical Review Team has the ability to gather additional telemetry from various points in the DNS to enhance their assessment of the name collision risks going on.

Which finally puts us to Stage 3. Then the Technical Review Team will submit the recommendation to the ICANN Board. The application may propose mitigation plans if there is a high risk for name collisions, which would repeat this process. But ultimately, the ICANN Board will make the final decision on proving the application and potentially assigning the string to a name collision list if the risks are so high for that particular string.

In summary, regarding data collection, we're not in 2012. You simply just cannot reuse the name collision detection mechanisms that we used back then. Controlled interruption back in the last round was effective to a certain extent, but it had limitations. It did not work for IPv6. Also, the name collision data was tightly tethered to the day in the life data collection from the root servers DNS, which now, due to DNS protocol evolutions and data collection mechanisms, that data is highly impaired compared to what it is, making the assessment for name collisions much more difficult.

With that, to seriously analyze name collisions, you have to collect the data in a different way. It's impossible based just on a general use case of what you observe at the root. That assessment only at the root may really impact the ability or the risk factor of particular names that are being assessed. With that, I think ICANN Org has expressed some concerns about risks in privacy and confidentiality. I think Ram is going to speak a little bit more about that, and I am happy to hand it back to him here.

RAM MOHAN: Thanks, Matt, and we'll ask that you hold your questions to the very end of this presentation. All of this that Matt presented was significant work. Oh, sorry, Matt, you have this slide, and then I think you hand it off to me.

MATT THOMAS: Oh, thanks, Ram. The only other thing that I would say around all of this is that clearly the Technical Review Team or the TRT is a central aspect of how to deal with name collisions going forward. It's a very technical function that needs to be fulfilled. It's a highly skilled activity. There's no generalized solution for name collisions.

One of the things the discussion group enumerated numerous times is that every name collision string is a little bit of its own special snowflake. There's no one-size-fits-all solution to name collisions. To ensure the TRT really is successful, it needs to be staffed with people with the right skills to be adaptive and to understand the nuances of the rules and the responsibility, and it also needs to be able to have the

ability to look at historical data as well as the data going forward, to understand potential impacts such as gaming, manipulations, changes in the DNS hierarchy, impairments to the DNS telemetry, and to be able to incorporate that into their overall DNS assessment of name collisions.

With that, it's really a highly specialized function. Getting the right folks to be able to serve on that is a critical step forward for doing a name collision assessment going forward. Ram, back over to you.

RAM MOHAN: Thank you. Matt and Susanne helped move the discussion group, which is a cross-community effort, to delivering this document which Matt just summarized. After that, what the SSAC did was it took that report and it crystallized that into a set of specific advice that we provided to the ICANN Board. I won't read through everything in here, but really what our findings were was, there's two kinds of risk. There is risk that you delegate and you have the high risk of exposing personal data, and there is the privacy risk associated with assessing name collusion.

Our recommendation was straightforward. We said to the ICANN Board, "Adopt the risk management framework that the Name Collision Project is proposing. Because this is not about technology, this is not about money, this is really about risk and making sure that as new TLDs get rolled out, that you first measure and then you mitigate the chances of a security instability problem.

We also suggested that if you look at prioritizing things, the ICANN Board ought to prioritize the mitigation of delegation risk versus the

mitigation of assessment risk. I'll speak more to that on the next slide. One of the things that came about in our analysis and in the discussion was ICANN Org came in with really useful and a really constructive dialogue.

What they said was, "Hey, if you ask us (ICANN) to go and assess all these names that are being applied for, there is inevitably privacy data that is going to come and then we're going to be forced to handle, manage, and evaluate all those risks. Especially when the discussion group has said the visual interruption method that Matt was talking about and the individual interruption with notification method, those are methods where there's a possibility and a likelihood that actual PII (Personally Identifiable Information) will come through.

They said, "Perhaps a better model might be, let's get all the applications to come in, it'll go through the normal evaluation process, but when it comes to the actual evaluation of collisions, we'll only have that be done after the names are given to the applicants -- whoever is the winner, if you will, of the TLD. We'll say, "Okay, you've got the TLD, we're granting it to you, but before we allow it to be delegated and work on the internet, you've got to run these tests. Then, from the results of those tests, we'll decide whether it should actually run or whether there should be mitigation, et cetera."

Of course, that's a possible viable method, but we think that all that's doing is passing the buck. If you're going to be getting PII information like that, you're going to get that either when you do it at the start or you do it at the end of this process. So you're not actually reducing the privacy risk. You could argue that outsourcing that work to all these

people, all these organizations that are going to be applying for TLDs, you're going to have the potential of an uneven implementation of it. In some ways, if there is a privacy breach or a privacy issue, the only way you're going to know it is after the fact.

Then, your mitigation is to go to that person or that organization, contractually ask them to do things. I mean, it feels much more cumbersome. Yes, it is more work for ICANN Org, yes, there is a liability and concerns like that, but what we're really saying is, listen, bite the bullet and do the work once, do the work at the start.

We also think there's some efficiency to it because if you're evaluating it and you realize the data shows that it's a high-risk string, you don't even go through any of the processes of contention sets and auctions and things like that. You just say, "Sorry, this is mail, we're going to put it aside." So we just think there's a common sense element to it, but there's also a practicality element to it.

The ICANN Board eventually has to make the call. We've been having really productive discussions with some of the committees there, education the technical and the risk committee there of the risk balance, and really sharing with them that you don't actually walk away from incurring the privacy risk. If you think there is a choice, and that the choice of doing the evaluation at delegation time is the better choice, you really ought to think very carefully about that. That really is what we've said. That's our question.

One of the things that people have been saying is, with the framework that the discussion group has proposed, if this is done, it can be gamed. You can manipulate the data. If you're another party that wants to stop

somebody else from getting a TLD, you can fairly straightforwardly gin up a machine to go and create a whole bunch of collisions, which can then create a problem. So we think that might be an issue. We've not proposed a solution to it, we're saying we recognize that that's an issue and a problem.

Matt, is there anything you wanted to add? Other SSAC members who were involved in this work, do you want to add something before we turn this over to a discussion?

MATT THOMAS: That was really good, Ram. Open the conversations and questions here.

JONATHAN ZUCK: One of the discussions we're engaged in with the GAC and with Org is about contention resolution, generally, and preventing the circus that took place in the 2012 round with the private auctions and things of that sort. One of the recommendations of that, coming from us, is to have the auction contention resolved as part of the application process, basically. That would even be before the name collision analysis that that contention would be resolved.

And having read your report and talked about this issue, it seemed that it would help to mitigate this issue and that you wouldn't know who you were in contention with and therefore wouldn't be incented to try and manipulate or gain someone's application. Does that logic make sense to you? That's where I went with it when I read the report, in the context of what we were trying to propose.

RAM MOHAN: Here's the way I look at it. Typically, ICANN publishes, on whatever the day is, all the TLDs that have been applied for.

JONATHAN ZUCK: This would be before that.

RAM MOHAN: No, the auction contention would be before that. But this gaming thing is if ICANN publishes the list and the collision analysis is going to happen sometime thereafter, there's a potential. Warren?

WARREN KUMARI: There are also some sets of things where you can game ahead of time. For example, I'm Pepsi, I can start sending queries now for .coke. Then that just means when Coke applies, there will be a bunch of existing stuff. Even if you don't know who you're going to be in a contention set for, you can figure out who your competitors are and stop poisoning their names. There's a bunch more extra tricky ways of doing it.

JONATHAN ZUCK: For sure, that's true. Our desire to try and deal with contention resolution differently won't address that problem, which is competition. Marketplace competition, not for the same name, but just to keep you from getting the name you want or something falls outside of our objective of trying to deal with contention.

WARREN KUMARI: That's true. But if you're in an industry, you can make some fairly good guesses about who is likely to be applying what string. Yes, they're not going to be perfect, but there are a bunch of things. It was fairly clear that a bunch of people were going to apply for .corp or .green. There are a bunch of names which are fairly clearly things that people are likely to apply for. So, yeah, you won't know for sure, but you can make some relatively good guesses. But, anyway. Jim?

JAMES GALVIN: I think what's important here is to not conflate too tightly separate issues. There is an issue of dealing with contention sets. There's an issue of dealing with auctions. There's an issue of dealing with name collisions. The name collision report was very careful to not actually say where in the sequence of events the due diligence should be done. Because there is much more analysis that has to be considered by ICANN or in the broader scheme of things.

The only thing the discussion group moves towards, and SSAC reinforces, is that the name collision due diligence should be done before the TLD is granted. Now, I'll offer the following perspective, this is my own personal perspective, about when name collisions should happen.

I think they have to happen before the contention set is resolved. And the reason why is because a part of the name collision process and analysis is, if the string is found to be a high-risk string, that might be useful information to those in the contention set. Because then they

may want to decide, "Gee, I don't really want this string." Or more importantly, once you discover that something's a high-risk string, an applicant has to have the opportunity to create a remediation or mitigation plan.

That information has to be upfront. I, at least, would feel like, as an applicant, I'd want to know that I have that to deal with, and that might affect how I feel about my place in the contention center and what happens. But that's just an impact and a consequence that ICANN has to evaluate along with all of the other relationships of the name collision analysis to the rest of the due diligence. It's interesting in that way. Thanks.

JUSTINE CHEW: Two points. One is just to address what Jim just explained. I think, from a subsequent procedure's policy recommendation point of view, name collisions, as Jim correctly pointed out, there is no recommendation in terms of what to be used. Or I should say that the recommendation basically says that the existing name collisions framework would apply unless the Board decides to use something else. It doesn't talk about where that process lies in the application throughout the delegation phase. Which is why the NCAP study too, if it's approved by the Board, then it's up to ICANN Org to determine how to do it. The challenge with auctions is that there is an existing policy recommendation that talks about ICANN auctions of last resort.

Whether there's a need to reverse that, if we're going to push for a different contention resolution, that remains to be seen. I don't know the answer to that, to be honest. But the point I wanted to make was

that I think the NCAP Study 2 is a pretty remarkable piece of work. I find it extremely useful in terms of guiding the procedure for dealing with the risk of name collisions going into the next round. I really, really appreciated the fact that SSAC had this NCAP discussion group that was open to non-SSAC members to participate in that process of generating the NCAP Study 1 and NCAP Study 2 reports. I hope that there will be other opportunities where other non-SSAC members can also participate in some of the work that SSAC does. Thank you.

RAM MOHAN: Thank you, Justine. That is certainly our intent to continue to invite members of the community on topics that are relevant. I see Hadia on the queue. I have about 10 minutes left, I think, in the session. If there's anybody else who would like to be in the queue, raise your hands in the chat now. I see Satish, Alan, and Gopal. If there's anybody else, please raise your hands. Otherwise, we'll draw a line there and then move on. So, Hadia.

HADIA ELMINIAWI: I was just going to say that when the name collision assessment happens, it is very important. As Jim pointed out, and of course this is very different than the contention resolution, however, if you do it after delegation and going through maybe contention resolution and then -- so you go through all this and then have it delegated and then do the assessment, what if it is a high-risk string? And in all cases, data manipulation could also happen at that point in time.

What would prevent it from happening after delegation when the applicant needs to go through the assessment? Still, this could happen. Also, as I recall, the NCAP study said if it's deemed a high-risk string, then the applicant can propose a mitigation framework, and that mitigation framework could be looked into and maybe the string goes forward. So there is a path forward that needs to be looked at before, I think.

RAM MOHAN: I agree, Hadia. Matt, there is a question in the chat. Would you like to take that? Will you please read the question out and then provide your response?

MATT THOMAS: Sure, thanks, Ram. I think the question is, is the wildcard-controlled interruption still in vogue with ICANN? Registry operator will implement controlled interruption for 90 days by inserting special records in its TLD zone file. To that end, I want to, at a high level, speak about the recommendations coming out of the NCAP Study 2, in that we did not try to be overly specific to the Technical Review Team, but to rely on their implementation advice on a string-by-string basis.

But to that end, the NCAP Study 2 considered four different technical implementations in ways that DNS telemetry data could be gathered in a more collective or holistic view to provide a more material and complete view and name collision risks. I will specifically say that is in vogue. I think the spirit of the recommendations is to allow for the Technical Review Team to just implement, at their will, the proper

mechanisms to collect the correct amount of data that they feel necessary to do a named collision assessment while the report itself outlines four of the probably best techniques that the discussion group at its point in time could enumerate and provide technical guidance on.

RAM MOHAN: Thank you, Matt. Appreciate it. Satish.

SATISH BABU: Thanks very much for the work on SAC0124, both in terms of the content as well as the process. In my personal opinion, assessing the privacy risk upfront seems to be a more sensible way to address this problem. My question is a clarification question. The proposed framework seems to be based on a temporary delegation during the initial delegation. But what happens post-initial delegation and post-redelegation? Does this list continue off name collisions? Thanks.

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

MATT THOMAS: That's a complicated question. There's obviously a balancing act here in terms of the number of changes that are applied to the root when we're delegating a new applied-for TLD. Traditionally, back in 2012, that delegation happened essentially after granting the proposed manner herein, it is introducing that before. We, again, are not being

overly prescriptive in our recommendations in the NCAP report in terms of how that is implemented.

We did have invited guests from IANA in terms of the concerns about root zone changes and scalability. Those were incorporated into the root. I think it's going to be on a case-by-case basis. I think the situations in which the name collision risks are elevated in the scenarios that you're probably alluding to are far from the norm, so it's most likely not a repeatable or material event that we should be talking about, but it's one to note in the corner cases.

RAM MOHAN: Alan, you have the last word.

ALAN GREENBERG: First of all, I want to compliment you on your restrained description of your reaction to the suggestion that the testing be done after the fact and by someone else. I was very impressed. But to follow on that, the concept of deferring it until it's almost delegated and then asking the group you're delegating it to do the testing themselves, certainly reduces ICANN's liability and risks. On the other hand, it assigns the responsibility of doing the testing to a group that has a very large vested interest in saying there's no problem, and potentially assigns the privacy risk to a group which may be very untrustworthy. We have no way of assessing it. There's just so many reasons not to do that.

RAM MOHAN: I'm smiling. I mean, this is what we think as well. You know, we're hopeful, and I'm cautiously optimistic that the folks we engaged with on the Board, they were very attentive, they had read all the material, they came prepared, they asked really good questions. They seemed to understand the nuance underneath all of this, not just a reflexive, "Let's first protect," type of thing. So I have hope.

Jim says that the question will be in front of the Board fairly soon, and that's part of the reason why we're also here sharing it. Because, as a specific ask, I think it'll be very helpful from the ALAC to brief your board representative as well and make sure that they are up to speed with the nuances of the issue here. Thank you.

JONATHAN ZUCK: I think that's everybody. Yes, it's always great to have the choir over so that we can preach to each other. We have a lot of common logic, I think, between us, if not sense. I, once again, very appreciate you coming and joining us and sharing this stuff. I think as a general rule, we've had some conversations privately about recommendations you've made that seem to have gone nowhere with the Board, and things like that, and we are very interested -- particularly when they have end-user implications -- to resurface those things and to be a part of reminding the Board that something's in front of them and that something should be done about it too. Because you guys do a lot of great work and none of it should end up in the bottom drawer. Thanks, again, for all that you do, and thanks for coming.

RAM MOHAN: Thank you very much. Thanks for inviting us, and hopefully we've given you evidence to do more of this.

JONATHAN ZUCK: This meeting is adjourned.

YESIM SAGLAM: And thank you to our wonderful interpreters and tech staff.

[END OF TRANSCRIPTION]